

AD-A168 552

SCIENCE AND TECHNOLOGY ANALYST'S ASSISTANT REVISION(U)

1/2


ADVANCED DECISION SYSTEMS MOUNTAIN VIEW CA


A S CROMATY ET AL. JAN 86 TR-1035-01 N00014-82-C-0602

UNCLASSIFIED

F/G 9/2

NL





AD-A168 552



12

TR-1035-01

SCIENCE AND TECHNOLOGY ANALYST'S ASSISTANT

Andrew S. Cromarty
Thomas L. Adams
James F. Cunningham
Carl J. Tollander

January 1986

Final Technical Report (Revised)

For Period 20 August 1982 - 1 January 1986

Contract No. N00014-82-C-0602

Approved for public release; distribution unlimited

Prepared for:

Dr. William Condell
Office of Naval Research
800 North Quincy Street
Arlington, Virginia 22217



003 FULL COPY

The views, opinions, and/or findings contained in this report are those of the author(s) and should not be construed as an official U.S. Government position, policy, or decision, unless so designated by other official documentation.

86 6 5 002

201 San Antonio Circle, Suite 286
Mountain View, California 94040-1289
415 941 3912 FAX 415 949 4029

Formerly Advanced Information & Decision Systems

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1. REPORT SECURITY CLASSIFICATION Unclassified			10. RESTRICTIVE MARKING AD-A168 552		
2. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited		
7. DECLASSIFICATION/DOWNGRADING SCHEDULE			5. MONITORING ORGANIZATION REPORT NUMBER(S)		
4. PERFORMING ORGANIZATION REPORT NUMBER(S) TR-1035-01			6. NAME OF MONITORING ORGANIZATION Office of Naval Research		
6a. NAME OF PERFORMING ORGANIZATION Advanced Decision Systems		6b. OFFICE SYMBOL (If applicable)	7a. ADDRESS (City, State and ZIP Code) 800 N. Quincy Street Arlington, VA 22217-5000		
6c. ADDRESS (City, State and ZIP Code) 201 San Antonio Circle, Suite 286 Mountain View, CA 94040-1270		7b. ADDRESS (City, State and ZIP Code)			
8a. NAME OF FUNDING/SPONSORING ORGANIZATION		8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-82-C-0602		
8c. ADDRESS (City, State and ZIP Code)		10. SOURCE OF FUNDING NOS.			
		PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT NO.
11. TITLE (Include Security Classification) Science and Technology Analyst's Assistant					
12. PERSONAL AUTHOR(S) Cromarty, Andrew S.; Adams, Thomas L.; Cunningham, James F.; Tollander, Carl J.					
13a. TYPE OF REPORT Final		13b. TIME COVERED FROM 82/08/20 TO 86/1		14. DATE OF REPORT (Yr. Mo., Day) January 1986	
15. PAGE COUNT 158					
16. SUPPLEMENTARY NOTATION					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB GR	expert systems; user interface; radar analysis; knowledge-based systems; distributed processing; remote database access; conflict resolution; signal analysis		
09	04				
15	04				
19. ABSTRACT (Continue on reverse if necessary and identify by block number) We describe a research prototype software system named ASTA (Assistant for Science and Technology Analysis) that employs artificial intelligence technology to interpret radar signals and infer the probable nature of the radar system that produces them. Such analyses are of interest, for example, to engineers who wish to "reverse-engineer" an electronic system. Architecturally, ASTA is a multiple-process consulting expert system operating in a personal workstation environment. It comprises two principal components: an expert system that reasons using knowledge of electronic system design, radar physics, and signal analysis techniques, and a multi-form-based user interface that maintains a separate database describing the complete state of interaction with the user, including a simplified user preference model, an interactive help subsystem, and a database of prior (completed or partial) analyses. These two processes are loosely coupled; that is, they operate in					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT. <input type="checkbox"/> OTHER USERS <input type="checkbox"/>			21. ABSTRACT SECURITY CLASSIFICATION Unclassified		
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. William Condell			22b. TELEPHONE NUMBER (Include Area Code) (202) 696-4672		22c. OFFICE SYMBOL

19. ABSTRACT

strictly separated namespaces, and they communicate using a mixture of synchronous (remote procedure call) and asynchronous (true) message passing. Message-based synchronization techniques are employed to assist in maintaining consistency between the two processes.

Major technological contributions from this project include:

- Development of a knowledge-based system that infers electronic system structure and specific design principles from a description of observed function and general design principles, rather than the other way around.
- Use of a simple mixed-initiative spreadsheet-style interaction metaphor to simplify the interface between the complex internals of the expert system and the possibly inexperienced user.
- Development and demonstration of specific techniques for detecting sources of conflict, both within an expert system and between the expert system and the external world, including a model of how such conflicts can be presented to and managed by the user.
- A proposed model of how expert systems can interact with remote databases, through interposition of intermediary data management agents and elevation of database queries to the status of active procedural entities.

A research prototype of the system we describe is operational. Unless explicitly noted otherwise, all functions described in this document are functional and have been demonstrated. A pair of actual processing scenarios are described in detail in the penultimate chapter of this report.

Accession For	
NTIS	<input checked="checked" type="checkbox"/>
DTIC	<input type="checkbox"/>
Unann	<input type="checkbox"/>
Just	<input type="checkbox"/>
By	
Dist	
Avail	
Dist	
A-1	



TABLE OF CONTENTS

	Page
1. INTRODUCTION AND OVERVIEW	1
2. STRUCTURE OF THE ASTA SYSTEM	4
2.1 KNOWLEDGE BASED SYSTEMS	4
2.2 ASTA ARCHITECTURE AND INTERPROCESS COMMUNICATIONS	8
2.2.1 ASTA's Multiple-Process Architecture	8
2.2.2 Maintaining consistency: The COP communication and control system	8
2.3 THE MRS INFERENCE ENGINE	12
2.4 EXPLANATION SYSTEM	16
2.5 CHANGING VALUES	18
2.6 PROCEDURAL ATTACHMENTS	20
2.7 CONVERSING WITH AN EXTERNAL DATABASE	22
2.8 SYSTEM DATABASE MATCHER	25
2.9 REASONING WITH UNCERTAIN INFORMATION	26
3. ASTA'S RADAR KNOWLEDGE	28
3.1 REPRESENTATION OF RADAR WAVEFORMS	29
3.2 TRANSMITTER ANALYSIS	33
3.3 ANTENNA ANALYSIS	34
3.4 PERFORMANCE ANALYSIS	40
3.4.1 Maximum Detection Range	40
3.4.2 Engagement Analysis	49
3.4.3 Unambiguous Range	51
3.4.4 Unambiguous Velocity	52
3.5 RESOLUTION AND TRACKING ACCURACY	52
3.6 SUBCLUTTER VISIBILITY MODELING	55
4. THE ASTA USER INTERFACE	56
4.1 OVERVIEW	56
4.2 SIMPLE PRESENTATION SURFACES	57
4.2.1 MPS icons	58
4.2.2 SPS components	58

4.2.3 Granularity of change reporting	59
4.3 MPS SUPPORT FOR SIMPLE PRESENTATION SURFACES	59
4.3.1 Display Engine	60
4.3.2 Cursor locator and positioning device	60
4.3.3 Object finder	61
4.3.4 Database of objects and object representations	61
4.3.5 Update relations between display representations and objects	61
4.3.6 Query and update by application control structures (EPS)	61
4.3.7 SPS activator and interpreter	62
4.4 CONSTRAINTS ON THE INTERFACE IMPOSED BY THE CURRENT IMPLEMENTATION	62
4.4.1 Display Engine	62
4.4.2 Inter-Process Communication	63
4.4.3 Defstruct data structures	63
4.4.4 C Language Interface	63
4.5 ASTA--MPS ARCHITECTURE OVERVIEW	64
4.6 EXTENDED PRESENTATION SURFACES (EPS)	65
4.7 REPORT GENERATION FACILITY	66
4.8 CHECKPOINTING FACILITY	66
4.9 FUTURE DIRECTIONS FOR THE USER INTERFACE	67
4.9.1 Consolidation of EPS interpreters	67
4.9.2 Subsystems	67
4.9.3 Help	68
4.9.4 SPS	68
4.9.5 Checkpointing	68
4.9.6 Constraint checking	69
4.9.7 Pop-up and pull-down menus	69
4.9.8 Presentation Graphics Support	69
4.9.9 Mouse support	69
5. CONFLICT DETECTION AND RESOLUTION	70
5.1 SOURCES OF CONFLICT	71
5.2 WHAT DEFINES A CONFLICT	72
5.3 RESOLUTION OF CONFLICTS	73
5.4 IMPLEMENTATION	74
5.5 EXAMPLES OF THE USE OF CONFLICT DETECTION	76
5.5.1 Reasoning From Multiple Points of View	77
5.5.2 Inconsistency Between Qualitative Abstractions and Numerical Calculations	78
5.5.3 Rule Inconsistency	79
5.5.4 Measurement or Assumption in Error	79
5.5.5 Indications of New Technology	80
6. SUBCLUTTER VISIBILITY MODELING	81
7. EXAMPLE PROCESSING SCENARIOS	91

**S&T Analyst's Assistant
CONTENTS**

Advanced Decision Systems

7.1 ANALYSIS SCENARIO 1: PULSED DOPPLER FIRE CONTROL RADAR	92
7.2 ANALYSIS SCENARIO 2: MECHANICALLY-SCANNED EARLY WARNING RADAR	121
8. SUMMARY	141
APPENDIX A. BIBLIOGRAPHY	144
APPENDIX B. ASTA KEYBOARD FUNCTIONS	146

LIST OF FIGURES

	Page
2-1: Structure of knowledge-based hypothesis formation systems	6
3-1: Pulsed-Doppler search and track radar waveform	31
3-2: Composite representation of radar waveform	32
3-3: Example transmitter output tube rules	35
3-4: Variation of antenna beam with smoothness of the illumination function	37
3-5: ASTA's antenna beam coefficients (rectangular apertures)	38
3-6: Antenna beam aperture analysis relationships (rectangular apertures)	39
3-7: Single-dwell cumulative detection and resolution/accuracy taxonomies	44
3-8: Inferring the Signal-to-Noise ratio for pulsed doppler radars	47
3-9: Anti-air warfare engagement diagram	50
6-1: Subclutter visibility vs. MTI, cancellers, and transmitter type	89
6-2: Subclutter visibility example	90
7-1: Pulsed Doppler scenario: Parameters for search frame f1	96
7-2: Pulsed Doppler scenario: Parameters for search frame f2	97
7-3: Pulsed Doppler scenario: Parameters for search frame f3	98
7-4: Pulsed Doppler scenario: Parameters for search dwell d1	100
7-5: Calculation of dwell unambiguous range and velocity for search dwell d1	101
7-6: Pulsed Doppler scenario: Transmitter form	102
7-7: Pulsed Doppler scenario: Radar System summary form after first ELINT collection	107
7-8: Pulsed Doppler scenario: Antenna form (uniform aperture illumination)	108
7-9: Pulsed Doppler scenario: Antenna form (cosine-squared aperture illumination)	110
7-10: Pulsed Doppler scenario: Radar Range equation (inferred terms)	111
7-11: Pulsed Doppler scenario: Radar range equation used to estimate peak transmitter power	112
7-12: Pulsed Doppler scenario: Results of using Radar Range equation to estimate peak transmitter power	113
7-13: Pulsed Doppler Scenario: Frame characteristics of PSK pulse-compressed tracking frame f4	115
7-14: Pulsed Doppler scenario: Resolution and measurement accuracy at the frame level	116
7-15: Pulsed Doppler scenario: Measurement accuracy at the dwell level	117
7-16: Pulsed Doppler scenario: Characteristics of search-track sequence	118

**S&T Analyst's Assistant
FIGURES**

Advanced Decision Systems

7-17: Pulsed Doppler scenario: Target form	119
7-18: Pulsed Doppler scenario: Interceptor form	120
7-19: Pulsed Doppler scenario: Radar System analysis summary	122
7-20: Pulsed Early-warning scenario: Search frame f1	125
7-21: Pulsed Early-warning scenario: Search frame f2	126
7-22: Comparison of frame/dwell definitions for (a) circular scan and (b) step scan	127
7-23: Pulsed Early-warning scenario: Search dwell d1	128
7-24: Defining sequence by antenna beam and scan characteristics	130
7-25: Pulsed Early-warning scenario: Sequence form	131
7-26: Pulsed Early-warning scenario: Antenna form	132
7-27: Pulsed Early-warning scenario: Transmitter subsystem inferences	133
7-28: Pulsed Early-warning scenario: Radar Range equation inferred terms	134
7-29: Pulsed Early-warning scenario: Maximum Detection Range inference results	136
7-30: Pulsed Early-warning scenario: Target characteristics	137
7-31: Pulsed Early-warning scenario: Interceptor analysis	138
7-32: Pulsed Early-warning scenario: Radar System summary	140
B-1: ASTA keypad layout	146

LIST OF TABLES

	Page
3-1: Swerling Model parameters for Neuvy's equations	46
3-2: Swerling Model numbers for typical collection contexts	48
6-1: Single canceller subclutter visibility coefficients	86
6-2: Double canceller subclutter visibility coefficients	87
6-3: Technological limits on subclutter visibility	87
6-4: Default values for clutter velocity standard deviation (m/s)	88
7-1: Logical inference structure for transmitter output tube type	103
7-2: Rules that lead to klystron as output tube type	105
7-3: Pulsed Early-warning scenario: Method of deriving radar range parameters	135

Acknowledgements

The authors wish to acknowledge the efforts of the many people who have contributed to the ASTA project, especially: Victor Askman, John L. Allen, Eric Domeshek, Ellen Drascher, Nancy English, Bruce Eskin, Milissa Feeney, Milton R. Grinberg, Sonia Lyris, Brian P. McCune, Joseph Mitola III, James Smith, David S. Spain, James L. Whitaker, Gerald A. Wilson, and Richard P. Wishner.

1. INTRODUCTION AND OVERVIEW

This report documents a two-year research and development effort into the design and construction of a prototype knowledge-based computer software environment for analyzing radar signals. Such analyses are of interest, for example, to engineers who wish to "reverse-engineer" an electronic system. That is, given a description of the system's observed behavior, they attempt to infer its structure and its functional capabilities.

Our contribution to the analysis process takes the form of a prototype computer-based consulting system named ASTA (Assistant for Science and Technology Analysis). The heart of ASTA is an "expert system" -- that is, a software system that incorporates the knowledge of radar domain experts and employs artificial intelligence techniques to reason using that knowledge. The expert system is embedded in a workstation environment that includes a high-quality and convenient user interface and mechanisms for transparently providing access to external databases of technical and historical data.

In addition to its computational and reasoning role in the analysis process, ASTA assists the analyst in two important ways. First, it maintains a database of partial or completed analyses for each ASTA user, and it permits those analyses to be modified, merged, or used as the basis for excursions into further hypothetical radar system analyses. Second, it provides a help and explanation facility that is useful to senior analysts and especially instructive to junior analysts. In particular, ASTA can explain in English what conclusions it drew from a given set of data, and why. This serves an important tutorial function for the novice, while providing expert analysts with a means of checking the integrity of the system's expert knowledge and reasoning process.

Numerous research topics in expert systems were addressed in the course of this project. Specifically, research was undertaken to address at least the following issues: detecting, reasoning about, and resolving conflict; reasoning under uncertainty; providing access to multiple heterogeneous external databases; effective user interface design for a mixed-initiative consulting system; integration of symbolic reasoning and numeric processing within a coherent computational framework; and design and implementation techniques for a distributed expert system. This report summarizes our principal research findings, our implementation experiences, and the status of the software system as of the publication date of this document. The remainder of this report is structured as follows:

Chapter Two is an Introduction and Overview to the ASTA system -- a description of its internal architecture, including the distributed systems technology that it incorporates; its inference engine; its support for explanation facilities; its methods for providing access to external databases; and some comments on its techniques for reasoning under conditions of uncertainty.

Chapter Three, "ASTA's Radar Knowledge," describes ASTA's expertise in radar physics and radar analysis. This discussion addresses both qualitative and quantitative modeling and reasoning techniques employed by ASTA to analyze radar systems.

A major component of the ASTA project, measured in both software and invested hours, is the User Interface subsystem that presents the state of the analysis to radar analysts and interacts with them during cooperative problem-solving. The fourth chapter, "The ASTA User Interface," describes that subsystem.

"Conflict Detection and Resolution" is the topic of Chapter Five. This chapter documents our research into techniques to extend a state-of-the-art inference engine by permitting it to explicitly represent and reason about conflict in its data, conflict in its rule base, and conflict between user-provided and inferred

hypotheses.

A topic that received special attention fairly late in the life of the ASTA project is subclutter visibility modeling. We describe our efforts in this problem domain in Chapter Six.

Chapter Seven, "Example Processing Scenarios," presents detailed examples of the use of ASTA in the solution of two realistic radar analysis problems, including discussion of the changes in the state of the inference engine and figures portraying the displays seen by the human analyst during the analysis session.

Our results and future directions are summarized in Chapter Eight.

2. STRUCTURE OF THE ASTA SYSTEM

In this chapter, we discuss the internal software architecture of the ASTA system. This discussion begins with an introduction to artificial intelligence (AI) software development techniques, and includes a treatment of distributed AI systems design as well as an overview of inference techniques employed by ASTA. We then discuss ASTA's explanation facilities and its support for pursuing alternative hypothetical lines of inquiry. Database topics are considered next, starting with a preliminary discussion of ASTA's internal mechanism for linking its reasoning system to the user interface's database and proceeding to a treatment of techniques for dealing with internal databases, external databases, and uncertain data.

2.1 KNOWLEDGE BASED SYSTEMS

AI systems usually perform complex inference that involves combining the use of a number of heuristics in an appropriate fashion to solve a problem. Virtually all applied AI systems are *knowledge-based*, in that they possess knowledge and techniques of a narrow but demanding application domain sufficient to exhibit a high level of expertise. When this knowledge has been derived from human experts, the system may also be called an *expert system*.

ASTA is tailored to radar systems by virtue of its database of symbolic facts and heuristics that define how new information should be inferred from existing evidence and previously reached conclusions. Because of the heuristic nature of much of the information that such a database contains, it is properly called a *knowledge base*. The ASTA knowledge base represents the first known effort to structure information about radar design as a function of observable operating characteristics, rather than from the point of view of the design

process. The knowledge comes from expert radar designers and analysts and from radar design handbooks. This includes general knowledge about radar systems, such as the physics of radar signals and the relationship between different components of radar systems. ASTA also has knowledge about itself: it contains explicit rules identifying (a) the ways it can use logic to solve problems, (b) the problems that are worth solving, and (c) the order in which interesting problems can most effectively be solved.

ASTA's knowledge may express either numeric or logical relationships. By separating declarative knowledge about radar physics and radar analysis problem-solving techniques from the generic inference-related control aspects of the computer program (its *inference engine*) that operates on that knowledge, ASTA facilitates inspection of its knowledge base and its line of reasoning in order to explain why new values were inferred. Further, its body of radar facts and analysis techniques can be modified by non-programmers, and it can operate robustly in the face of the partial or errorful data that typify analysis tasks.

Figure 2-1 depicts the basic structure of knowledge-based hypothesis formation systems. All such systems have the four parts shown: they reason about input data to determine the best hypothesis (or hypotheses), using a knowledge base of facts and heuristics about the application domain. Although many hypothesis formation tasks deal with sensor data, the same system structure applies to purely symbolic tasks such as textual information retrieval. The reasoning process is usually carried out by an "inference engine" that is capable of reasoning about facts and heuristics regardless of the particular subject area with which they are concerned.

Knowledge-based systems have a number of important attributes that make them attractive for automating analysis tasks:

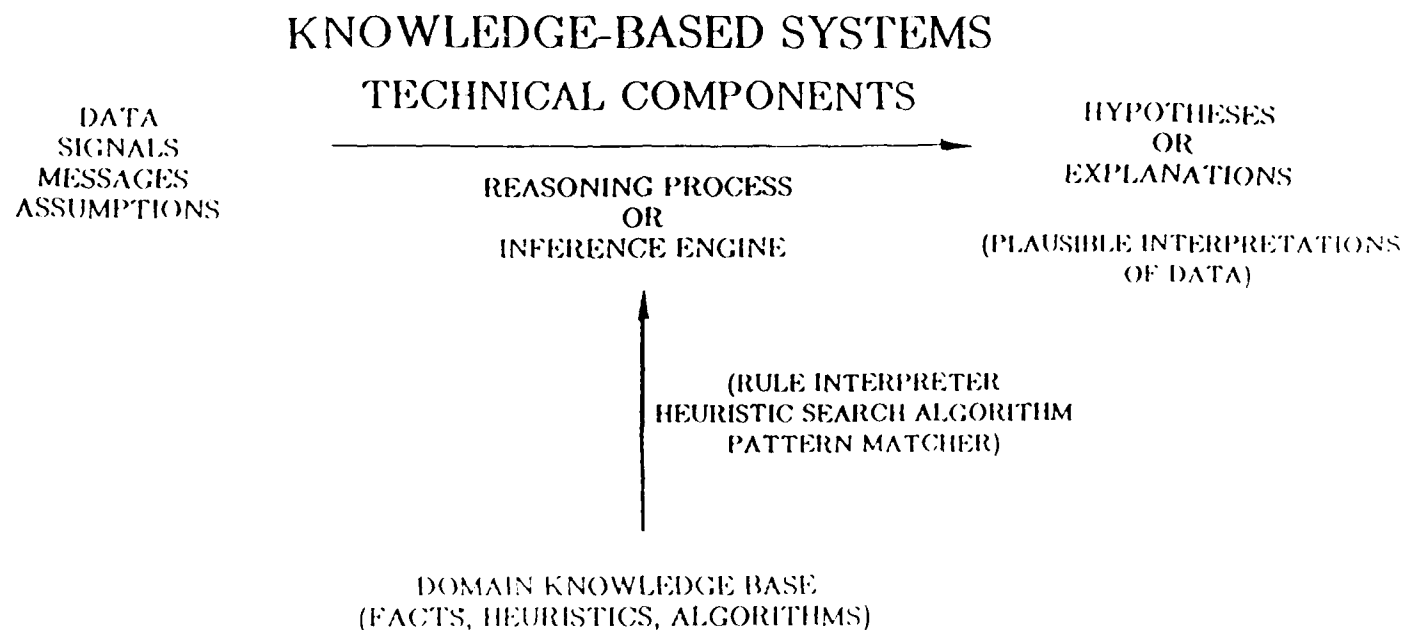


Figure 2-1: Structure of knowledge-based hypothesis formation systems

- They make more analysis possible by replicating the expertise of the (presumably scarce) human experts.
- They can carry out analyses in locations where human experts do not exist or do not care to go.
- In contrast with humans, knowledge-based systems are good at handling the myriad details of complex situations.
- In contrast with other computational approaches that are more formal and algorithmic, knowledge-based systems are more robust: they are designed to deal with problems exhibiting uncertainty, ambiguity, inaccuracy, and missing data.
- The behavior of knowledge-based systems, by virtue of their relatively independent and nonprocedural knowledge bases, is more easily and efficiently modified and may be more readily understood by users who are nonprogrammers.
- Knowledge-based systems may also be more efficient than other automation techniques, by virtue of their ability to first determine what aspects of the current problem are critical and to then devote most of the computational resources to solving them.

Knowledge may express either numeric or logical relationships. By separating this declarative knowledge from the procedural computer program that operates on it, ASTA can explain why new values were inferred, can have its long-term operations modified by a non-programmer, and can operate robustly in the face of missing data.

2.2 ASTA ARCHITECTURE AND INTERPROCESS COMMUNICATIONS

2.2.1 ASTA's Multiple-Process Architecture

The ASTA system is designed as a multiple-process system, specifically with one process for the user interface and a second process for the inference engine. The tasks performed by the user interface include providing a presentation surface for the user (the ASTA forms) and the screen control primitives such as *select form*, *display form*, *move cursor*, and so forth. The interface also provides its own database describing the design of the forms and the present state of those forms with respect to the values and flags associated with each system attribute. The inference engine provides its own separate but consistent database as well as a database of rules and a computational model for manipulating the system attributes. Both of these processes will be discussed in more detail in further sections. The two processes must maintain a consistent database of attributes and communicate when their state changes using the COP communication and control system.

2.2.2 Maintaining consistency: The COP communication and control system

The inference and user interface processes contain data that are largely distinct semantically: rules are under the purview of the inference process, while the details of the presentation surface are the domain of the user interface alone. The two processes do, however, experience overlap in several semantic domains, most notably for (a) data entered by the user and (b) inference results that are to be presented to the user. Because multiple copies of these data exist in the two

name spaces and because of their time-varying nature, the potential exists for inconsistency between the two databases. For example, the inference process could erroneously work on obsolete data if newly entered data held in the user interface database are not yet installed in the inference process's knowledge base; similarly, incorrect results (or none at all) could be displayed to the user even after the inference process had determined the proper values, if the corresponding entries in the user interface database have not been updated.

We need not solve all the problems addressed by a distributed DBMS such as SDD-1 [5], since only one data source (the user or expert system) can be updating the database at any moment in time. ASTA is, however, faced with the problems of maintaining multiple copies of a single logical database that a distributed DBMS must solve. In order to ensure consistency and integrity of the common data across the two processes, ASTA employs a *time-stamping* data communication architecture with message routing between the two processes explicitly managed by a semi-intelligent controller. The time-stamping, message-passing architecture from which ASTA has been constructed is the COP system [2].

In its full generality, the COP communications and control system provides for resource management and planning functions as well as communications; within the ASTA implementation, however, we use it primarily to ensure timely and reliable interprocess communications. COP provides to each of the two client processes (user interface and inference) a message-passing view of the external computational environment. All communication between modules is effected by means of a single **send-message** primitive having the following structure:

(send-message from to timestamp class [text])

where *from* and *to* identify the sender and receiver, *class* is a member of the set of permissible remote operations which the sender may invoke (in the case of ASTA, the set {**SSet**, **SGet**, **Show**}), and *timestamp* is a structured message

identifier that includes a timestamp timed according to the sender's clock and guaranteed to be unique across all clients in the system. Where messages take arguments, *text* specifies those arguments, in keyword (attribute-value) form. (The *from* value is provided by the communications slave, described below, to minimize the potential for implementation errors on the part of the client module's designer and to prevent message forgery.)

Each client is transparently provided with a communications slave that performs two functions:

- It dispatches messages sent by the client using **send-message**.
- It responds to incoming messages, either directly or by dispatching them to the appropriate client function.

If an incoming message is a member of the client's set of permissible messages, message receipt is transformed into a call to the client's corresponding function. If the message is not a member of the client's operator set, the slave determines whether it is an otherwise known message class (such as a bookkeeping request from the controller which the slave itself can execute); if not, an error message is transmitted back to the controller by the slave.

As described above, our approach to ensuring consistency for inferred facts and hypotheses relies upon *procedural attachments* within the inference process. These procedural attachments are essentially explicit per-rule and per-relation specifications of the corresponding relation(s) in the companion database that depend upon the value of the attached rule or relation. For example, the inference engine is instructed to check for and execute procedural attachments whenever a new fact is inferred in order to ensure immediate update of the user

interface database.

Depending upon the specifics of the rule or relation so attached, the procedural attachment may specify either a *remote procedure call* form of update, wherein the inference engine waits until the companion database has been updated before proceeding, or a pure *message delivery update*, in which messages are queued for delivery to the companion process but the sender does not wait for delivery. (In the latter case a remote procedure call is guaranteed to wait until all pending update messages have been delivered and acknowledged.)

Emulation of remote procedure call is supported by allowing the sender to "block" itself until a return message is received; this blocking action is a primitive of the communications slave, which continues to listen to the input port and process incoming messages, either queueing them on an agenda of tasks to feed the client once it is resumed or executing them directly if they are communications-specific (such as a request for a message indicating the client's status). This restricted use of message passing permits two clients to synchronize their state: for example, the sender client may ask for a datum from the target client's database and then wait, blocking all other analysis- and display-related operations within its name space until a (time-stamped) return value is received.

In the current ASTA prototype system, a highly simplified version of the COP design has been employed. Because the agents in this system communicate using UNIX "pipes," they do not suffer from many of the difficult problems that more demanding COP applications have faced. ASTA's present interprocess communication implementation is currently used primarily to provide a simple asynchronous and synchronous remote procedure call capability between the two processes.

2.3 THE MRS INFERENCE ENGINE

The knowledge base of facts, rules, and metarules is contained entirely within a process that is assigned the responsibility of pursuing the inferences required to support the analysis. MRS provides the general purpose inference engine and data base capability used within ASTA to store and derive radar system parameters. It supports a rule-based approach that employs a knowledge base of facts and rules along with a flexible control structure used to guide the inference strategy. MRS is a domain-independent reasoning and representation system in which knowledge about any field may be represented. In the ASTA system, MRS stores and maintains all of the domain specific knowledge of radar systems including the initial default values for physical constants, the current known radar parameters that have been entered by the analyst, the radar parameters that have been derived from one or more known parameters, the rules used to relate the radar parameters to each other, and the meta-level knowledge used to control the use of the rules and data. The information that the system uses may be numeric or symbolic in nature. MRS has the ability to make mathematical calculations or to draw inferences based on symbolic information to derive new symbolic information. Furthermore, symbolic information can be used to select the appropriate form of a calculation or to provide constraints on the range of values in the terms of an equation. Because the analysis domain is characterized by information that is informal, imprecise and incomplete, we store all of this knowledge declaratively, rather than procedurally. By storing the inferring procedures declaratively, the system has the ability to reason with them, manipulate them, and use them only when enough information exists to derive new data from the existing data. The inference process is therefore very flexible and can use the information it has available to make all the conclusions it can, but will not be hindered or rendered useless when certain radar parameters are unknown.

The **stash** operation is used to store *assertions* (facts and rules describing the current problem domain) in the database. These assertions are stored as n-tuples in MRS. Intermediate inferences need not always be stored in the database, so the meta-level control is used to define when and where assertions are stashed. For the ASTA application, we have chosen to represent the radar system parameters in 3-tuples consisting of a property, object, and value, where the object can be thought of as an index into a table of values for all objects with the specified property. For example, the 3-tuple (**prf f1 15**) states that the property "prf" has the value "15" for the radar frame indexed by "f1." Assertions are retrieved from the database by the **truep** ("truth" predicate) operation. It queries the database for evidence concerning the validity of a statement given the current context. The system will determine if this statement is true by searching the database for it; if the fact is not present, MRS will then try to infer the validity of the statement from the database using the known facts along with the rules of inference.

The rules in MRS are of the common *if-then* production rule form, with the *if* clause consisting of one or more antecedents (preconditions) that, when true, imply that the consequent statement associated with the *then* clause is true. Any logical proposition can be encoded as a set of rules in this form by first putting it into conjunctive normal form and converting each disjunct into an appropriate rule. MRS then allows these rules to be used in two ways: either in a data-driven, forward chaining direction, from antecedents to consequents, or in a goal-directed, backward chaining direction, from consequents to antecedents. In the data-driven direction, the system will try to match the information it has against the antecedents of the rules, and when successful, will add the consequent of the matched rules to the database. Conversely, when performing goal-directed reasoning, the system will hypothesize that a particular goal is true and try to find rules that contain the desired goal in its consequent. If such rules are found, the system will attempt to match the antecedents of these rules against the database, and if at least one rule is successful, will then add the desired consequent to

the database. If no such rule succeeds, the system may try to find further rules with the unsuccessfully matched antecedents in the consequent of another rule and attempt to determine the truth of these using the same procedure. For example, ASTA has the following rule in its database:

```
(if (and (pulse-modulation $f PSK)
         (chip-duration $f $p))
    (compressed-pulse-duration $f $p))
```

The antecedent assertions are the clauses with the properties "pulse-modulation" and "chip-duration," while the consequent clause has the property "compressed-pulse-duration." (Variables are distinguished by a "\$" preceding the name, and can match against any instantiated term in another clause.) In this example, the first antecedent has the variable "\$f" which will match the first index it finds with the property "pulse-modulation" that has the value of "PSK." When using this rule in the forward chaining direction, the system will try to match these antecedents, which must both be true since they are joined by the "and" operator, against the database. If successful, it will assert the consequent with the property "compressed-pulse-duration" using the same instances for the variables that it used for the antecedents. As an example of applying the previous rule in the backward chaining direction, suppose it is desired to determine the compressed-pulse-duration of a particular frame "f1" as the instantiation for "\$f." The compressed-pulse-duration can be thought of as a goal that the system will try to prove using the database of current assertions along with the applicable rules. MRS does this by instantiating the antecedents of this rule; if these can be matched against the database, MRS will use the same instantiation in the consequent clause to assert the fact. If one or more of the antecedents does not occur in the database, the system may post them as new goals and try to prove them true using the same backward chaining mechanism.

The method of search and chaining direction is selected by the ASTA design team on a case-by-case basis and implemented through the use of the meta-level rules (with default search techniques employed where appropriate). Meta-level control operations provide control over the use of the rules as well as a means of manipulation of the environment of assertions that are currently valid. The meta-level consists of both assertions and rules of the same form as the base level rules and assertions, stored in the same database. The meta-level assertions dictate how the system uses rules. For example, they can proscribe the use of backward chaining or forward chaining mechanisms, control the stashing of results in the database, or specify whether it is appropriate to seek out more than one instantiation of a particular goal.

Meta-level assertions can be used to affect the current context of applicable databases that are valid. MRS allows multiple databases, called *theories*, to be used for stashing both rules and assertions. The meta-level defines which theories are currently "active" (searchable) and how to change the state of a theory between active and inactive. The meta-level rules may contain context dependent conditions that determine when the meta-level assertions are applied. These rules have the same form as the base level rules; however, the base level rules embody knowledge about the particular domain, whereas the meta-rules adapt the use of those rules to the current situation. This mechanism allows the system to understand its effect in the current context and then adapt to the constantly changing situations that it encounters.

These database mechanisms allows ASTA to work with multiple, mutually exclusive hypotheses simultaneously. Competing hypothesis data are manipulated in different databases, and storage and retrieval are managed using meta-level rules. Specifically, the ASTA system will activate a theory to store the hypothesized values and values derived from inferences made using these hypotheses. The theory may also contain specific rules that are only activated with the theory and, therefore, only with the one set of data associated with the

particular hypothesis. If the hypothesis is found to be valid, the assertions made in its local theory may be moved to a more global theory, or this theory may simply be moved into a new context within the global framework. If the hypothesis is found to be invalid, the assertions may be discarded or moved to an inactive state which will not affect the other theories.

The ASTA knowledge base also contains a network of *justifications* that are used to generate explanations. Every assertion in the system has an associated explanation of how it was derived, either from the user as an input parameter, from a system constant, or a value derived from a rule. For each assertion that is derived from a rule, the explanation database must save the rule that was used to derive the fact, the current meta-level control that was in effect when the rule was fired, and a binding list of the instantiations of the variables in the rule. From this information, the system can trace the derivation of a value, and the rules that were applied to the input parameters to infer the value that is in question.

2.4 EXPLANATION SYSTEM

One of the most important capabilities of expert systems is their ability to provide explanations for the conclusions they reach. Expert system explanation mechanisms allow a user to follow the reasoning path backwards from the conclusions that were reached, through the inference engine, back to the original primitives, either input assertions given to the system by the user or physical constants agreed upon as legitimate by the domain community. ASTA has this capability to trace inferences based on its own specific inference engine and primitive system values.

The ASTA explanation system is deeply embedded in the MRS inference engine. Internally, MRS solves problems by unifying facts in the database with a template, an n -tuple formally called a *clause* representing the problem to be

solved. The system initially looks for an appropriate fact in the database and returns this as the solution to the problem, but if no fact is present, the inference engine finds rules with a consequent that matches the problem template and tries to find a solution for each antecedent of this new rule. This process continues until all of the antecedents have been matched, or otherwise, no solution is reached. This is the backward chaining inferencing mechanism described in the previous subsection. MRS also provides forward chaining inference in which facts are immediately matched against rule antecedents: if all or a rule's antecedents are present in the database, its conclusion is asserted. With this type of reasoning, no specific query has to be asked by the user; the system tries to infer as many facts from the data as it can. The explanation facility must therefore have the ability to trace inferences drawn in both backward chaining and forward chaining directions. Due to the homogeneous representation of facts and rules used for both forward and backward chaining, MRS uses the same set of rules and facts for both inference mechanisms, and it records the inference technique used to derive a specific fact and the supporting evidence together with the fact itself. The explanation facility can therefore access all of the relevant information about the derivation of a fact from a standard structure.

Similarly, MRS allows information to be stored about rules within the system. This information includes an English-like explanation of the rule and the template for the rule itself. The explanation, written in the jargon of the domain rather than in terms of the computer and system implementation, provides a bridge from one fact in the database to other facts. This explanation is entered into the system by the rule base implementor and is only modified when the rule base is modified. For facts (as opposed to rules), *template rules* are attached to the facts that specify how they can be converted into English-like phrases; when the fact must be presented as a part of an explanation, these template rules are employed to provide (create) human-readable text for the fact.

The ASTA explanation facility provides two types of explanations, the shallow *why* and the deep *explain all*. The *why* traces the derivation of an MRS fact only one level back in the chain of rules used to infer a fact. If the fact was entered by the user or is a system constant, the explanation facility tells you the source of the value. Otherwise, the value was derived through the use of a rule. In this case, the explanation facility provides three pieces of information. The first is the inference technique used to derive that specific value (either forward chaining or backward chaining). The second piece of information is a set of rule antecedents and the corresponding values used to derive the consequent fact. The third is the English language explanation for the rule used in the derivation.

In the *explain all* facility, the interrogated fact is traced back recursively through all of the rules that were used to derive it. If the fact was input by the user or is a system constant, the explanation is given and the rule tracing stops at this point. If the fact was derived through the use of a rule, again the same information is provided by the system as with the *why* facility. In contrast with the *why* facility, if the antecedents of this rule were derived through the use of a further rule, *explain all* would trace this rule in a similar manner until we have traced all of the derivations of the original fact back to the primitives from which it is derived.

2.5 CHANGING VALUES

The ASTA system allows the user to change values of system attributes during an analysis session; ASTA then updates the values of dependent attributes to reflect these changes. For example, the system has these rules in the rule base (we include the textual explanation as well as the rule-language form of the rules):

```
(if (and (type radar-antenna planar)
        (shape radar-antenna rect)
        (height radar-antenna $h)
        (width radar-antenna $w)
        (* $h $w $a))
    (area radar-antenna $a))
```

"The area of a rectangular radar antenna = height * width of the antenna."

```
(if (and (type radar-antenna planar)
        (shape radar-antenna hex)
        (height radar-antenna $h)
        (* .866 (* $h $h)) $a))
    (area radar-antenna $a))
```

"The area of a hexagonal radar antenna = (.866 (height ** 2))."

```
(if (and (area radar-antenna $a)
        (nominal-wavelength radar-xmtr $w)
        (gain-factor radar-antenna $b)
        (* (/ 12.566371 (expt $w 2))
           (* $a $b)) $j)
    (db-conversion $j $g))
    (gain radar-antenna $g))
```

"The gain of a radar antenna =

$$((4.0 * \text{PI}) * (a * g)) / (w ** 2)$$

where w = nominal wavelength of the radar system

a = the antenna area

g = a gain factor specifying the ratio of effective area to physical area."

We start with the assumption that all of the necessary attribute values have been entered and the gain of the radar antenna has been inferred based on these values. Specifically, the values for shape of the radar antenna is *rect*, and

the area has been inferred based on this value. Furthermore, the gain of the radar antenna has been inferred based on this value of antenna area. Now, if the user decides the rectangular radar antenna shape assumption is incorrect and changes the shape from *rect* to *hex*, the area of the antenna and the gain of the antenna will both be inconsistent. The ASTA system has a meta-level inferencing structure that will respond to this type of chaining in the database and, based on this new radar antenna shape value, will update all of the values that have been inferred based on the old radar antenna shape. In this example, both the area of the radar antenna and the gain of the radar antenna, which was based on the original radar antenna shape, must be changed.

Internally, the procedure for updating the database starts with a trace of the justifications used to derive values. If a value has been inferred based on this now changed value, the system has to remove the inferred value from the database. If this removed value was itself used to support the inference of other values, they now have to be removed as well. Once the entire chain of reasoning has been explored and inconsistent values removed from the database, the system will try to re-infer these attribute values based on the new value of the original fact that has been changed by the user. For each fact that was removed, the system initiates an inference operation to determine whether a value for the attribute of current interest is still inferable with respect to the new set of facts in the database. If so, the new value is added to the database and the system is again in a consistent state; otherwise this attribute is not inferable given the new state of the database.

2.6 PROCEDURAL ATTACHMENTS

Procedural attachments allow the inference engine to compute values directly, by providing an extension to the narrow theoretical framework of logic based systems. We will first provide a description of procedural attachments and how they are used in the ASTA system, followed by an explanation of (a) the

relationship of procedural attachments to logic oriented rule based systems and (b) why procedural attachments are necessary.

A procedural attachment is a function called by the inference engine to determine a value and/or produce a side effect when the system is asked to instantiate a specific clause template using the inference engine (i.e. determine a value of some system attribute). A meta-level assertion is the fact which tells the system that a certain clausal form is to be called as a function, rather than being looked up in the database of assertions as a fact or as the consequent of a rule to be solved using inference.

When MRS is asked to infer a radar system attribute, it must first lookup the inference procedure in the meta level assertions and perform the appropriate inference procedure. The typical inference procedures are either backward chaining or forward chaining. However, this third type of inference procedure, procedural attachment, may be used to perform a function type of inference. In MRS, which is written in LISP, the procedural attachments may be any of the standard LISP functions available with LISP or any function written by the user within the environment provided by the LISP implementation. For the ASTA system, the environment allows the use of LISP functions, custom written LISP routines, or even subroutines written in C, the language which is used as the basis of the UNIX operating system. MRS knows that the procedural attachment is appropriate because the system designer has (directly or indirectly) placed an assertion in the set of MRS meta-level inference procedure assertions.

In ASTA, we have taken advantage of all three types of procedural attachments. For example, we have used standard LISP procedures to perform some of the basic mathematical functions, such as addition, subtraction, multiplication, etc. We have also written some specific procedural attachments to perform basic operations such as conversions from dB to power domain or visa versa, determining multiplier factors for PRI and PRF, conversions from wavelength to frequency, and so forth. C routines are used to perform low-level machine bit

operations when necessary; for example, the truncation of low order bits from floating point numbers is achieved using a C procedure embedded in LISP to overcome an implementation flaw in the underlying FranzLisp language.

Let us consider an example, one wherein LISP functions are used to calculate the value of an arithmetic function. MRS uses first order predicate calculus as its basic theoretical framework. As a consequence, its inference procedures are based on *rules of inference* that can be applied to clauses to form new clauses, formally called *theorems*. The sequence of rule applications is called a *proof* of the theorem. The *rules of inference* tell us how we can use one or more facts, or clauses, to produce new facts. *Modus Ponens* is an example of a *rule of inference*. In order to determine the value of mathematical expression, such as "(* 5 3 \$x)," the database must have either the explicitly stated fact or a set of rules which can be applied to determine the value of "\$x" using a *rule of inference* and some facts stored in the database about basic mathematical operations. Obviously, it would be much simpler and more efficient to call a LISP procedure to evaluate this expression. Therefore, we have defined a special technique to exit the logic-based inference engine, execute some external LISP code which returns a value, and continue with the inference. This allows us to escape to a function outside of MRS when it is more convenient or efficient to evaluate the "truth of a clause" (such as an arithmetic relation) rather than write a set of rules to determine the value.

2.7 CONVERSING WITH AN EXTERNAL DATABASE

A wide variety of data bases are already in use by radar system designers and analysts, and clearly ASTA will benefit if it can capitalize on the prior existence of massive collections of relevant domain data. Unfortunately, these data bases are in general mutually incompatible; they are frequently poorly organized, and they are supported on a variety of DBMS's. An important task in the design of ASTA has therefore been the development of a uniform method of

providing the inference process with access to external database systems.

Our solution to this practical problem has been to isolate the knowledge about specific external database query languages and schemas in a single additional *database access expert* process that maps data requests from the inference engine to the appropriate database query. This prevents the inference process's knowledge base from becoming cluttered with arcane knowledge of the external DBMS, provides a potential degree of parallelism (in that inference can proceed while the database access expert formulates a database query), and modularizes the DBMS-specific knowledge so that no changes to the inference process's knowledge base are necessary (at least in principle) in order to support access to additional external databases or to change the query protocol as external databases evolve.

While this meets our immediate pragmatic goal of providing access to external DBMS's, several problems still present themselves during the construction of the database access expert. For example, a separate access capability (whether in a monolithic access expert or multiple independent such experts) must be provided for each external DBMS. Furthermore, the schema of the external DBMS must be duplicated in the database access expert, and thus changes in the external database schema still require a corresponding alteration of the database access expert. Finally, the database access expert must contain knowledge not only of the query language syntax, but of the computational semantics of query language constructs as well, in order to effectively map inference process requests into database queries. Whereas this semantic mapping is normally performed either by a human user of a conventional database or by a transaction designer, the database access expert cannot appeal to either source of interpretive expertise and thus must carry the additional burden of maintaining and applying that knowledge of query semantics itself.

On the basis of our experience with the expert system-external database interface, we observe that a more effective solution is at hand if data base systems provide an additional query capability not currently supported: that of *query predicate* satisfaction. That is, an ideal external database from the standpoint of the expert system would be one that accepts not merely a static query but a predicate that can be executed within the name space of the database.

Such a query predicate could contain, for example, a weighted vector of values or ranges which would be applied to candidate tuples in the database by the DBMS to produce a *degree-of-match* measure, where a database tuple is considered to itself be a vector in n -space. In the simplest case, attributes of a relation would be objects with considerable mathematical structure on them (such as values in \mathbf{R}^n with the elements of the basis set being semantically compatible -- for example, a database that contained only latitude and longitude information), and the metric for degree of match is little more than a variance calculation. In more difficult cases where, for instance, the range of data values for an attribute is nominal-level (i.e. no ordering relation applies), the query predicate must contain more information in order to convey the *goal* of the "user" (in this case, the inference engine) in posing the query. For instance, if the attribute in question were a spectral color, the degree to which "yellow" is a satisfactory match for a query that requested tuples that are "like orange" would be a function of the intent of the inference engine in using that color information, which intent must be reflected in the query predicate's handling of the color attribute. This predicate would then allow the database to perform a better-informed search (i.e. provide better recall performance) by virtue of semantically-derived information provided as a part of the query by the agent (the inference engine) formulating that query.

An obvious extension of this approach, and one especially useful in expert system applications, would be a DBMS that provides not only the tuples that match above some threshold or according to semantic constraints embodied in

the query predicate, but also the degree-of-match measure as defined by the query function for each tuple satisfying the query predicate.

A database system that supported query predicates would offer a potential performance improvement in several respects: raw query-satisfaction speed, decreased inter-process traffic, and decreased requirements for post-processing of query results by the inference engine, all by virtue of the knowledge contained in the query that focuses the database search process. Such cooperation between the expert system and the external database would greatly facilitate the development of knowledge-based problem solvers in practical problem domains. We know of no database systems that currently provide such a facility.

2.8 SYSTEM DATABASE MATCHER

We have examined the issues described in the section above and have build a system database matcher using these ideas. The database matcher uses the input and derived information from the present ASTA analysis and tries to match the analysis results with known attributes of present operational systems. Due to the imprecision of the data supplied to ASTA, an exact match to a known system would be very unlikely, and perhaps incorrect. Therefore, ASTA's matching system will perform a *fuzzy match* with the known radar systems, returning the five best matches from the set of radar systems in the database.

A capability of the matching system is the ability to explain major differences present between the analysis results and the known system results. This can be useful for either of two reasons. First, the analyst may wish to check the input values and assumptions, and perhaps when these assumptions were made on tenuous grounds, change the values using the change value facility and retry the system matcher. This may lead to a better analysis and a better system match. A second reason for displaying the differences between the matched system and the analysis is that the real system may have changed from the time the

system database was built, and the analysis reflects these new capabilities or features of the system. The system has alerted the analyst to the new features which can then be examined in more detail.

The system matcher uses five attributes to determine the identification of the analyzed system, a lower RF (carrier frequency) bound, an upper RF bound, operational modes of the radar system, maximum detection range, and the nationality of the system. The details of the matching and justification processing are discussed in further detail in another section.

2.9 REASONING WITH UNCERTAIN INFORMATION

The ASTA system has a general mechanism for handling uncertainty with respect to a discrete uncertainty space. The uncertainty space is defined by three discrete states: *{unlikely, likely, conflict}*. The uncertainty mechanism is very general and can be used for any set of system attributes that should be captured using uncertain reasoning. The method employed for uncertainty calculation is to check for all evidence confirming a fact and all evidence refuting a fact. If all of the evidence points to one attribute's presence or absence, then the uncertainty value will be likely or unlikely, respectively. If the evidence is pulling heavily towards both the supporting and refuting conclusions, then *conflicting* is asserted. The conflict state occurs when there is good evidence for two mutually exclusive facts, and therefore one of the results must be incorrect, due to either an incorrect input or a set of inconsistent rules.

For the uncertainty mechanism to be active, an uncertainty rule set must be defined for a particular set of domain rules, which can be easily completed using the general uncertainty mechanism. This mechanism uses the the MRS functions **provable** and **unprovable**. These functions query the knowledge base about the validity of a clause in the knowledge base. If it can be inferred from the given set of facts, then it is provable; if it cannot, then it is unprovable.

These functions are then used to query the database as part of a rule antecedent, and if the queries are satisfied, the consequent is asserted. For example, the transmitter tube type can be determined using this rule in its proof:

```
(if (and (provable (tube-type radar-xmtr $x))
         (unprovable (not (tube-type radar-xmtr $x))))
    ($x radar-xmtr likely))
```

This rule says if there is supporting evidence for a particular transmitter tube type and no evidence refuting it, then that tube type is *likely*. (The transmitter tube type will be bound to the quantified universally variable "\$x.") Obvious in this mechanism is the importance of the domain knowledge captured in the supporting evidence rules and the refuting (e.g. **not**) evidence rules. Similar rules are used for the other cases of uncertainty. This type of rule exemplifies one aspect of the general approach to nonmonotonicity ASTA uses to reason about uncertain evidence. Further details on how ASTA reasons using knowledge of radar systems appears in the next chapter.

3. ASTA'S RADAR KNOWLEDGE

The primary goal of the initial ASTA implementation was to demonstrate the feasibility of applying artificial intelligence techniques to interpret the information from a wide variety of data sources and to form self-consistent hypotheses about the interrelationships between the radar subsystems and associated weapons systems. It was not intended that the initial prototype develop a complete radar analysis package, but rather that it illustrate how a wide variety of techniques could be synergistically combined. As a result, the level of detail we provide in any particular analysis may not be extensive.

This chapter summarizes the radar knowledge that is currently included in the knowledge base. The range of applicability of the techniques described in this chapter is illustrated in Chapter Seven, by applying this knowledge base to two widely different radar systems.

This chapter describes ASTA's radar knowledge as the system currently exists. Since ASTA does not have complete knowledge of radar systems (e.g., complete knowledge of radar physics, historical trends, etc.), this chapter does not in itself constitute a thorough introduction to radar systems. Readers interested in a tutorial introduction to radar systems are referred to [6] and [7] for more detail.

We begin by describing the method chosen to represent the radar waveform symbolically. Then the rules for analyzing the transmitter and antenna subsystems are presented. The receiver is not analyzed as a separate subsystem, but certain parameters of the receiver that are crucial to performance analyses, such as the estimation of upper and lower bounds on the range tracking loop bandwidth, are considered. We then show how the waveform characteristics and subsystem inferences can be used to obtain predictions about the performance characteristics

of the radar, such as its detection and tracking characteristics and its target handling capacity. The methods for associating the radar system with a weapons system are then given. Finally, the evaluation of subclutter visibility is described.

3.1 REPRESENTATION OF RADAR WAVEFORMS

In a modern multifunction radar, it is common practice to employ a relatively small set of primitive waveforms in flexible patterns in order to achieve various operational objectives. For example, several pulse trains, each at different Pulse Repetition Frequencies (PRFs), may be transmitted to one location in order to disambiguate the range of a target, or multiple beams may be created by providing feed signals at different frequencies to a common antenna. This functional grouping of basic pulse structures leads naturally to a composite characterization of the waveform in which higher order pulse structures are defined in terms of combinations of the primitive pulse structures.


The pulse is the basic unit of characterization of a pulsed waveform. (We observe in passing that in waveforms exhibiting complex discrete intrapulse modulation, characterization at the subpulse level may be desirable.) Examples of the attributes of a pulse are its frequency, duration, and intrapulse modulation characteristics. Pulses may be grouped into various types of pulse structures, such as pulse bursts and pulse trains. Examples of the attributes of a pulse group are the duty cycle, and the number of pulses in the group and, in the case of uniformly spaced pulses, the pulse repetition frequency, or PRF.

The pulse structures may be further organized into commonly repeated units which we shall term *frames*. A frame is defined whenever a set of pulse groups is used repetitively to achieve a common purpose. For example, a pulse train and a modulated pulse burst may be used together to achieve a tracking and command guidance frame. Examples of the properties of a frame are its


duration, operational mode, and PRF mode.

In the current ASTA prototype, the lowest level signal structure is the frame; properties normally defined by radar analysts at the pulse and pulse group level and inherited by the frame are defined directly at the frame level. Thus pulse attributes, such as amplitude modulation and pulse group attributes, such as *pulse repetition intervals*, are defined directly at the frame level. This choice was made for the prototype because the inheritance rules for attributes at higher levels of the hierarchy are more interesting and because such a description was adequate to encompass the original class of signals under consideration -- multiple, discrete PRF, pulse doppler waveforms, and conventional pulsed waveforms with PRF stagger. We observe in passing that the latter can be viewed in many important respects as a special case of the former, wherein the number of pulses in the pulse train is equal to one, the number of staggered PRF intervals is equal to the number of discrete PRFs and the frame duration is equal to the interpulse interval.

The higher levels of the radar waveform hierarchy result from the operational characteristics of the radar. Figure 3-1 illustrates the waveform of an electronically scanned pulse doppler radar that is performing an autonomous radar fence search and simultaneously tracking targets handed off to it by dedicated search and acquisition radars providing high altitude search for the same weapons platform. The same basic waveform, consisting of the repeated transmission of two frames, is delivered to each spatial position prior to scanning the beam. The search pattern is broken periodically by a differently structured track dwell that is pointed toward the detected target. The periodicities of this waveform induce two higher levels of description on the waveform: the *dwell level*, which describes how the individual frames are combined to achieve a particular function, and a *sequence level*, which describes the mode of multiplexing the dwell functional units. The amplitude pattern of the search dwells reflects the antenna gain of the intercept antenna as the radar beam sweeps across it. These changes are slow



(a) Radar waveform (time domain)



(b) Antenna scan pattern (spatial domain)

Figure 3-1: Pulsed-Doppler search and track radar waveform




Figure 3-2: Composite representation of radar waveform

if the radar is scanning adjacent spatial beams positions and the intercept receiver is pointed in a fixed direction. A large change in amplitude can occur during a track dwell when the beam is pointed toward a target in track, since a different portion of the intercept beam is sampled. Figure 3-2 shows how the waveform of Figure 3-1 can be described as the composition of the waveform components described above. Sequences associated with each of the operational modes of the radar may be defined. ASTA may thus use intermediate results derived at the frame level and propagate these results to higher levels of the pulse-structure taxonomy to obtain efficient computation of the operational characteristics of the radar when multiple operational modes employing common pulse substructures are employed. For example, the unambiguous range of the dwell of a discrete multiple PRF system may be obtained by applying the Chinese Remainder Theorem to the PRFs of its constituent frames and the tracking accuracy of a dwell whose frames are noncoherently combined can be obtained by applying the conventional improvement formula (the inverse of the variance equals the sum of the inverses of the component variances) to the accuracy of its constituent frames.

The description of waveform compositionally provides not only computational efficiency but an understanding of the operation of the radar over time, since once the operational modes have been isolated and classified, they can be used to segment the received waveform and thus establish the operational sequence employed by a radar during a given exercise.

3.2 TRANSMITTER ANALYSIS

The primary goals of the transmitter analysis are determination of the oscillator configuration and output tube type. Determination of oscillator configuration is based primarily on PRF analysis. The combination of high PRF and high duty cycle is taken as evidence of pulsed doppler signal processing, which requires a coherent oscillator configuration. Pulsed systems that do not require coherent processing techniques are assumed to use a power oscillator.

Evidence for and against each major category of tube type (klystron, cross-field amplifier, travelling-wave-tube, magnetron) is accumulated based on bandwidth requirements, power requirements and operating frequency. If there is evidence supporting the selection of a particular tube type and no evidence against it, then that tube type is assumed to be a likely candidate. Conversely, if there is no evidence supporting the use of a particular tube type and there is evidence suggesting it should not be used, then the output tube type is designated as unlikely. If there is evidence both in favor of use of a particular tube type and also suggesting it should not be used, then the conflicting evidence is indicated for the appropriate tube type.

Figure 3-3 illustrates some typical rules used in transmitter output tube type inferences. The first rule expresses the fact that if both supporting evidence can be found and no refuting evidence can be found, then a particular tube type should be considered likely. The second rule, an example of refuting evidence, states the bandwidth limitations of klystrons. The third rule, which might be defined as a rule of minimum complexity, expresses the fact that a magnetron is often used for incoherent, pulsed applications and should be considered likely, unless there is contrary evidence in the database.

3.3 ANTENNA ANALYSIS

The aperture illumination (distribution of energy across the aperture) and the far-field pattern are Fourier transform pairs. In many cases, it may be possible to obtain approximations to the antenna dimensions, thereby defining the extent of the aperture, but the amplitude distribution across the aperture may be unknown. However, the analyst may still gain considerable insight into the gain and measurement capabilities of the radar by observing the effects on the beam characteristics of interactively manipulating a symbolic amplitude distribution smoothness parameter. The symbolic amplitude distribution parameter is used because high accuracy is not required and this approach facilitates table lookup.

```
(stash-rule%explanation
  '(if (and (provable      (tube-type radar-xmtr $x))
            (unprovable(not (tube-type radar-xmtr $x))))
    ($x radar-xmtr likely))
"If there is supporting evidence for a particular transmitter tube type and
no evidence refuting it, then that tube type is likely.")

(stash-rule%explanation
  '(if (and (frac-inst-band radar-xmtr $x)
            (> $x 0.12)
            (coherency radar-xmtr yes))
    (not (tube-type radar-xmtr klystr)))
"If the instantaneous bandwidth of a coherent transmitter is more than 12%,
then it is probably not a klystron.")

(stash-rule%explanation
  '(if (and (modulation-type pulsed)
            (coherency radar-xmtr no)
            (unknown (not (tube-type radar-xmtr magntr)))))
    (tube-type radar-xmtr magntr))
"A pulsed noncoherent transmitter is likely to be a magnetron unless there is
evidence in the data base refuting it.")
```

Figure 3-3: Example transmitter output tube rules

Instead of specifying the amplitude distribution analytically, calculating its Fourier transform, and then searching the resulting function for the locations and amplitudes of local maxima and minima, the analyst is provided with a menu of increasingly smooth symbolic descriptors that can be entered in the aperture illumination slot on the antenna.

Figure 3-4 shows the variation in the shape of the antenna beam as the illumination function is varied from sharp (uniform) to smooth (cosine). The smoother illumination function results in a decrease in gain/efficiency, broadening of the main beam and a decrease in sidelobe energy. The template used by ASTA to determine the beam coefficients for rectangular apertures is shown in Figure 3-5. The illumination function/beam coefficient trends illustrated pictorially in Figure 3-4 are shown diagrammatically in Figure 3-5.

The nominal frequency obtained from ELINT observations and the antenna dimensions can be used in the traditional formulas (see Figure 3-6) to obtain first order estimates of the gain and beamwidth. The gain and beamwidth modification factors of common aperture types are stored as a function of symbolic aperture smoothness parameter in a rule format (see Figure 3-5). The gain factor shown in figure 3-5 is for line sources. The total gain factor to be used in the gain expression of Figure 3-6 for rectangular antennnas is the product of the gain factors associated with the horizontal and vertical dimensions. If the analyst does not specify a smoothness parameter, then the uniform distribution is assumed. The primary tuning parameter is usually the antenna first sidelobe level. For example, if a radar fence search mode is used, then sidelobe levels must be low to avoid excessive ground clutter. The default value of uniform aperture illumination gives an unacceptably high sidelobe level, so the analyst selects the cosine-squared illumination, which results in a much reduced sidelobe level at the expense of a broader beam and less gain. This example illustrates one way that ASTA can be used to support analyst guided refinement of first-order calculations when insufficient data is available for a complete analysis.





Figure 3-4: Variation of antenna beam with smoothness of the illumination function

Illumination Function		Gain Factor	Beamwidth Multiplication Factor		First Sidelobe Level
Uniform		1.0	60.8		-13.2
SMOOTHER ↓	Cosine	0.81	69		-23
	Cosine Squared	0.67	83		-31

↑ EFFICIENCY ↓ MAIN BEAM BROADENING ↑ SIDELOBE ENERGY

Figure 3-5: ASTA's antenna beam coefficients (rectangular apertures)

BEAMWIDTH

$$\Theta = (\text{BEAMWIDTH-MULTIPLICATION FACTOR}) \times (\lambda/d)$$

GAIN

$$G = \frac{4\pi A}{\lambda^2} (\text{GAIN FACTOR})$$

Figure 3-6: Antenna beam aperture analysis relationships (rectangular apertures)

3.4 PERFORMANCE ANALYSIS

In this section, the techniques employed by ASTA for determining the performance characteristics of the radar are described. Analysis of the detection capabilities of the radar is presented first. Then the determination of resolution and tracking accuracy is discussed.

3.4.1 Maximum Detection Range

There are three different approaches to estimating maximum detection range: the *radar range equation*, *engagement analysis*, and *unambiguous range analysis*. ASTA computes maximum detection range by each of the methods whenever the required input data are available. If the results are not consistent, a conflict is reported.

The radar range equation relates the transmitter power, antenna gain, transmission frequency, system losses, signal processing characteristics (such as integration time), target cross section, signal to noise ratio required for detection, and detection range of the system in a single expression. If inferences are available for $n-1$ of the terms in the radar range equation, then the n th term may be derived. The purpose of engagement analysis is to relate the target and interceptor kinematics to the detection range required for the radar. Unambiguous range is the maximum range at which a target's position can be determined unambiguously. Images of the target at multiples of the PRF expressed in distance units (scaled by the speed of light) appear at distances greater than the unambiguous range. Each of these concepts is discussed in greater detail in the following sections.

3.4.1.1 Radar Range Equation

ASTA uses one basic radar range equation to solve any problem concerning detection range in a noise limited environment. The appropriate values for the terms in the equation are selected according to the appropriate context. Coherent integration over the entire pulse train is assumed for pulsed doppler systems in the current implementation. The radar range equation can be written as:

$$P = \frac{(4\pi)^3 \cdot R^4 \cdot (S/N) \cdot (kT) \cdot N_F \cdot L}{G^2 \cdot \lambda^2 \cdot \sigma \cdot T_D}$$

where

P = transmitter power

T_D = dwell time on the target

R = radar target range

S/N = signal-to-noise ratio

k = Boltzmann's constant

T = noise temperature

N_F = noise figure

L = system losses

G = antenna gain

λ = nominal wavelength

σ = target radar cross-section

For pulsed systems, P in the above radar range equation is the peak transmitter power and T_D is the inverse of the signal bandwidth or equivalent pulse duration. For pulsed doppler systems, P in the above radar range equation is the average transmitter power and T_D is the inverse of the Doppler filter

bandwidth or the time on target.

3.4.1.2 Determination of Required Signal-to-Noise Ratio

The radar system designer usually specifies detection performance objectives in terms of probability of detection and probability of false alarm. The radar range equation requires the specification of detection performance in terms of the signal-to-noise ratio required to achieve the specified performance. This section discusses the relationship between these two views of detection.

The signal-to-noise ratio is the ratio of signal power at the time of peak signal output to the average noise power at the receiver terminals. The signal to noise ratio can be enhanced by both coherent and non-coherent integration. Coherent integration improves signal-to-noise ratio by a factor that is roughly proportional to the number of pulses integrated for small signal-to-noise ratios, but it requires more complex circuitry in order to maintain adequate pulse coherence. The time interval over which coherence can be maintained sets an upper bound on the coherent integration time. Non-coherent integration is easier to accomplish but enhances signal-to-noise ratio at a rate that is proportional only to the square root of the number of pulses integrated.

Detection characteristics also depend on the amplitude fading characteristics of the target. If the time scale of target fading is much shorter than the time interval between detections, then the target amplitude statistics of adjacent pulses are essentially independent. If the time scale of target fading is much greater than the time interval between detections, then the target amplitude statistics of adjacent pulses are highly correlated. In most situations, detection performance is superior for uncorrelated fading statistics than it is for the correlated fading case. The detection modeling used by ASTA takes both integration time, fading amplitude statistics and temporal correlation into account in computing the required signal-to-noise ratio. In the current implementation,

integration over the pulse train in a pulsed doppler frame is assumed to be coherent. Interpulse integration in a pulsed system is assumed noncoherent. The relationship between detection performance at the frame and dwell levels is calculated approximately by Nuevy's equation, which is an analytic relationship specifying the relationship between the probability of detection for a single dwell, the false alarm probability and the number of pulses (pulse trains) integrated noncoherently in the dwell. The relationship between detection at the dwell and sequence level is based on the expression for cumulative probability of detection for a number of independent dwell detections of known probability of false alarm and detection probability. The number of independent dwell detections is based on inferences of the time interval available for detection and the search revisit interval.

Figure 3-7 shows schematically the detection probability relations. The frame SNR is determined from the pulse (train) characteristics and the effective radiated power. The dwell signal-to-noise ratio is related to the frame signal-to-noise ratios via Nuevy's equation using assumptions about target fading characteristics and the single dwell detection probability. The dwell detection probability and the cumulative probability of detection are related by inferences about the number of opportunities for detection and assumptions about the track initiation scheme employed.

In the following paragraphs, we first discuss the use of Nuevy's Equation to relate frame and dwell detection characteristics. Then we illustrate how the dwell detection characteristics can be combined with the cumulative probability of detection and search operational heuristics to obtain a relationship between dwell detection characteristics and overall detection characteristics.

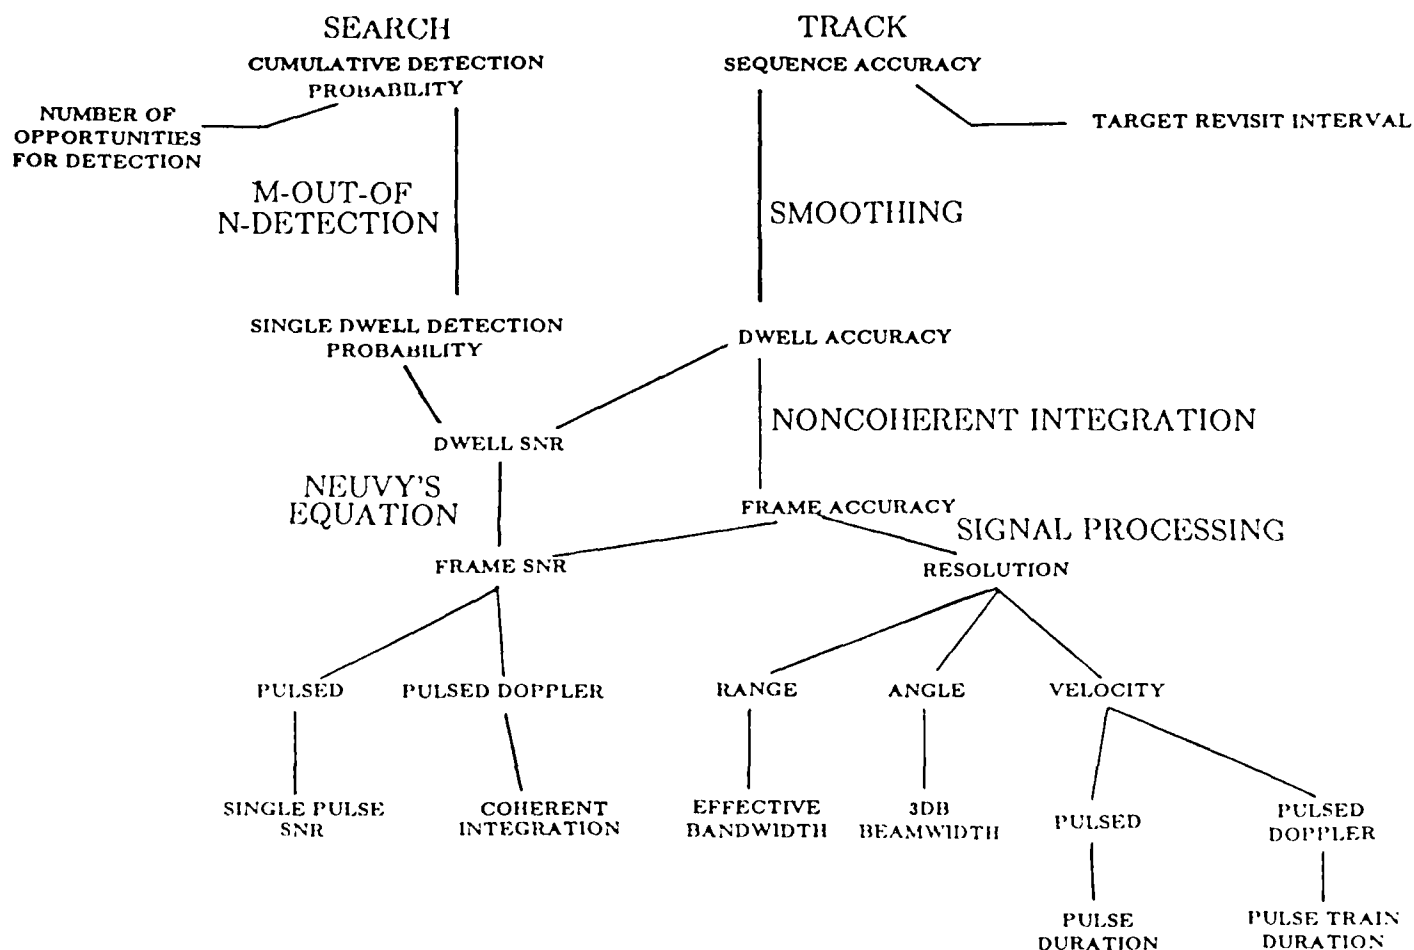


Figure 3-7: Single-dwell cumulative detection and resolution accuracy taxonomies

3.4.1.3 Nuevy's Equation

A closed form expression relating the required signal-to-noise ratio in the pulse (train), S/N, and the desired detection/false alarm characteristics for Swerling target fading models has been obtained by Nuevy. These expressions are accurate to about +/- 1dB for ranges of detection probability/false alarm rates of interest. The probability of detection is given by:

$$\log_{10} P_D = - \left[\frac{\alpha \log_{10} n_f}{(S/N) n^\gamma} \right]^{\frac{1}{\beta}}$$

where

P_D = the probability of detection per dwell

P_{FA} = the probability of false alarm

n_f = the false alarm number = $0.693/P_{FA}$

n = the number of frames integrated noncoherently per dwell.

$\alpha = y_1(y_2 - a \cdot y_3)$

$a = e^{\frac{-n}{3}}$

and y_1, y_2, y_3, β and γ are parameters given in Table 3-1.

In Table 3-1, the Swerling Model Numbers refer to the following assumptions about target amplitude fading:

Swerling Case 1: Constant power exponential.

Rayleigh amplitude, exponential power. Echo pulses received from a target on any one dwell are of constant amplitude throughout the entire dwell.

Swerling Case 2: Independently fading exponential.

Target cross-section probability density function the same as for Swerling Case 1.

Table 3-1: Swerling Model parameters for Neuvy's equations

Swerling Parameter	Case 1	Case 2	Case 3	Case 4	Case 5
y_1	$\frac{2}{3}$	$\frac{2}{3}(1 + \frac{n-1}{2n})$	1	1	1
y_2	1	1	$\frac{3}{4}$	1	1
y_3	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{1}{2}$	0	$\frac{2}{3}$
β	1	$0.275 + \frac{0.725}{n}$	$\frac{2}{3}$	$\frac{1+4a}{6}$	0.275
γ	$\frac{2}{3}$	0.64	0.64	0.6	0.56

Amplitudes are assumed independent from frame to frame.


Swerling Case 3: Constant power Chi-square.

Target power probability density is Chi-square with four degrees of freedom.
Constant amplitude from frame to frame.

Swerling Case 4: Independently fading Chi-square.

Target cross-section probability density function the same as for Swerling Case 3.
Amplitudes fade independently from frame to frame.

Swerling Case 5: Nonfluctuating Target.



TIME INTERVAL FOR DETECTION = 15.8 SEC SEARCH-REVISIT-INTERVAL = $50 \times (3 \times 5.8 \text{ msec}) = 0.87\text{s}$

NUMBER OF DWELLS COMBINED FOR DETECTION = $\lceil 15.8\text{sec} / 0.87\text{sec} \rceil = 18$

CUMULATIVE PROBABILITY OF DETECTION = 0.9

PROBABILITY OF FALSE ALARM = 10^{-6}

$$P_D^{CUM} = 1 - (1 - P_D^{DWELL})^N \quad (\text{AT LEAST ONE DETECTION})$$

PROBABILITY OF DETECTION PER DWELL = 0.22

NEUVY'S EQUATION (SWERLING CASE 1, 3 FRAMES PER DWELL)

SNR REQUIRED = 5.5dB

Figure 3-8: Inferring the Signal-to-Noise ratio for pulsed doppler radars

Table 3-2: Swerling Model numbers for typical collection contexts

Swerling Model	Scatterer Distribution	Frame-to-frame Fluctuation	Dwell-to-dwell Fluctuation
1	Rayleigh	Constant	Incoherent
2	Rayleigh	Incoherent	--
3	One Large and Many Small Scatterers	Constant	Incoherent
4	One Large and Many Small Scatterers	Incoherent	--

Table 3-2 relates the Swerling model numbers to typical collection contexts. Swerling cases 1 and 2 result from a large number of independent scatterers. The distribution associated with Swerling cases 3 and 4 is representative of a large reflector together with other small reflectors or one large reflector subject to small changes in orientation.

Figure 3-8 illustrates how the SNR required for detection may be inferred. Starting with an assumption about the radio designer's desires for cumulative detection probability and a track initiation criterion (at least one out of N detections chosen here for simplicity) the required detection probability on a dwell can be determined, if we can establish the number of opportunities for detection N . The latter can be established for a pulsed doppler system by locating the largest uneclipsed range interval near the maximum detection range as shown in Figure 3-8. This interval can be converted to a time interval for detection by scaling by the nominal target velocity. The number of dwells combined for detection is obtained as the ratio of the time interval for detection and the search revisit interval. The latter is obtained as the duration of the search sequence, assumed

in this example to consist of 50 dwells, that is, the search volume scanned by the antenna consists of 50 spatial beam positions that are cyclically revisited. Once the required dwell detection probability has been obtained in this way, the frame required can be calculated using Nuevy's equation, the number of frames in the dwell, some target fading assumptions and an assumption about the desired false alarm rate.

3.4.2 Engagement Analysis

The interrelation of the key elements of a surface-to-air missile system are examined very clearly and conveniently in the anti-air warfare diagram. ASTA implements the single-shot, point defense (self defense) form of the engagement diagram, but extension to area defense and consideration of multiple shoot firing doctrines is possible. As shown in Figure 3-9, time-to-go is plotted versus range. Time-to-go means the time remaining until the target (aircraft or anti-ship missile) arrives at the defended ship. A constant velocity is assumed for convenience for both interceptor and target, although acceleration is easily accommodated in these diagrams. The straight line marked target has a positive slope inversely proportional to its velocity. The interceptor starts from zero range and flies towards the target. Its slope is negative and inversely proportional to its velocity. The interceptor line intersects the target line at the maximum effective range of the interceptor. The maximum detection range is determined as the sum of the effective range of the interceptor and the distance traveled by the target in the sum of the flyout and reaction times. The flyout time is determined from the interceptor maximum velocity and maximum effective range. Reaction time is the elapsed time interval between target detection and interceptor launch, and includes the time required to establish track, identify the target and assign a missile battery to the target.




Figure 3-9: Anti-air warfare engagement diagram

3.4.3 Unambiguous Range

If the time required for the pulse to reach the target and return is greater than the interpulse spacing, then an ambiguity arises as to whether the return was obtained by reflection from a close-in-target from the previous pulse or by reflection from a target at a greater range illuminated by an earlier pulse. The maximum range at which range can be determined unambiguously is called the unambiguous range of the radar.

For simple pulsed radars the unambiguous range, R_u can be determined directly from the PRF, f_p , using

$$R_u = \frac{c}{2 f_p}$$

where c is the speed of light.

The unambiguous range can be increased by using several fixed PRFs, measuring the ambiguous range in each PRF and comparing the measurements to eliminate ambiguities. The PRFs are chosen to have a common submultiple frequency and are usually related by the ratios of closely spaced relatively prime integers. ASTA contains an algorithms for estimating the unambiguous range from the PRF information of its constituent frames.

The algorithm first calculates a ranging PRF using a statistical harmonic analysis procedural attachment. The ranging PRF is the PRF that, if used in a single frame, would result in the same ambiguous range as that of the dwell. The unambiguous range expression given above can thus be applied to the ranging-PRF to obtain the unambiguous range of the dwell.

The statistical harmonic analysis procedural attachment finds a set of integer multipliers that have no common submultiples and, when multiplied by the corresponding PRF, result in the same ranging PRF. The common methods of least common multiple analysis cannot be exploited because the measured

PRFs are not exact and because the PRFs may contain offsets equal to one or two clock cycles that destroy the exact integral relationships among the PRFs. The procedure used attempts to find a multiplier that results in an approximately integral value of the ratios of the PRFs. The algorithm attempts to find a set of multipliers under three different assumptions about the degree of fit and then accepts the largest multiplier obtained and multiplies this value by the minimum PRF to obtain the ranging PRF.

3.4.4 Unambiguous Velocity

The unambiguous velocity is the largest velocity at which it is possible to measure velocity unambiguously. For simple pulsed radars the unambiguous velocity V_u , is proportional to the ratio of wavelength, λ , to the PRF, f_p .

$$V_u = \frac{1}{2} \frac{\lambda}{f_p}$$

The unambiguous velocity of a multiple PRF system can be found by calculating a fundamental prf for the dwell. The fundamental PRF is obtained by determining the least common integer multiple of the PRFs. The least common integer multiple is obtained by applying the statistical harmonic analysis procedure to the PRFs comprising the dwell.

3.5 RESOLUTION AND TRACKING ACCURACY

Resolution refers to the ability of the radar to separate two or more targets that are close together, and to determine in range, angle or doppler frequency (radial velocity) that there are two targets instead of one. Accuracy is expressed as the square root of the mean-square error, i.e., rms error. In the current implementation of ASTA the accuracy formulas represent the accuracy of parameter

measurement when there is no other object near the target to affect the measurement.

The angular resolution, $\Delta\theta$, is assumed to be proportional to the half-power beamwidth, θ_o ;

$$\Delta\theta = K_R \theta_o$$

where the resolution factor K_R , ranges from 1 to 2 (1.5 typical).

The range resolution ΔR , is assumed to be proportional to (scaled by the half-speed of light) the equivalent half power pulse width, τ_e ,

$$\Delta R = K_R \left(\frac{c}{2} \right) \tau_e$$

The equivalent half power pulse width is taken as the inverse of the signal bandwidth so that spread spectrum range resolution improvement techniques, such as phase-shift keying (PSK) and linear FM, are properly accounted for.

The velocity resolution, ΔV , is assumed to be proportional to the ratio of wavelength to the time on target, T_f ,

$$\Delta V = \frac{K_R}{2} \frac{\lambda}{T_f}$$

where for pulsed systems

T_f = the pulse length

and for pulsed doppler systems

T_f = the length of the pulse train.

Measurement accuracies are reported at each level of the waveform hierarchy. In the current implementation only errors due to thermal noise are reported. The frame measurement accuracy is taken as proportional to the resolution, with a constant of proportionality that depends on the type of signal processing employed, and inversely proportional to the square root of the signal-to-noise ratio.

For angle measurements the constant of proportionality is

$$\frac{1}{\sqrt{2}} \frac{K_R}{K_M}$$

where K_M is the normalized monopulse slope, which ranges from 1.0 to 2.0 (1.89 typical). For range measurements, the constant of proportionality is

$$\frac{1}{2 \sqrt{2} K_R}$$

For velocity measurements, there is no accuracy reported for pulsed systems at the frame level and for pulsed doppler systems the constant of proportionality is

$$\frac{1}{\sqrt{2} K_R}$$

The inverse of the variance of the measurement error at the dwell level is taken as the inverse of the sums of the inverses of the variances of the measurement errors of the constituent frames. For pulsed systems, the velocity measurement standard duration at the dwell level, σ_{T_v} , is inversely proportional to the range measurement standard deviation, σ_{T_R} , at the frame level, with the constant of proportionality related to the number of frames in the dwell, n , and the time interval between frames, T_s , by

$$\sigma_{T_v} = \frac{2 \sigma_{T_R}}{n T_s}$$

At the sequence level, the dwell inverse variances are summed, as at the dwell level, to obtain the sequence inverse variance, but tracking bandwidth limits are applied so that the number of dwells smoothed never exceeds the bounds imposed by target acceleration and jerk characteristics.

3.8 SUBCLUTTER VISIBILITY MODELING

ASTA contains a limited ability to reason about subclutter visibility, using knowledge about both technological limits to clutter processing and fundamental radar physics (such as details of clutter spectrum spread). For a more detailed discussion of ASTA's treatment of subclutter visibility, see Chapter Six.

4. THE ASTA USER INTERFACE

In this chapter we describe the User Interface of the ASTA system, that is, the portion of the ASTA analysis system responsible for interaction with a human analyst.

4.1 OVERVIEW

The ASTA user interface makes use of the concept of "surfaces" as multi-dimensional abstractions that handle interaction between the user and the ASTA application program. This is similar to the idea of a "blackboard" to which both the user and an application program may write and from which both may read.

The "reading" and "writing" is done by adding, manipulating and deleting the graphical or textual "objects" on these surfaces which represent further declarative and procedural abstractions to the user and to the application.

A user interface may maintain a number of these surfaces, or *windows*, for a given application, along with the means for an application and the application's users to write and read from them. Thus we see that one problem of user interfaces is not to get the user to interact with the application, but rather to get the user and the application to interact with an ordered set of intermediate surface structures.

Most choices regarding user interface architectures are those of how and when surfaces and constellations of surfaces may be manipulated. We shall describe the means that we have developed to deal with these choices.

The ASTA user interface combines the set of tools and techniques for maintaining these surfaces as provided for in the Multipurpose Presentation System

(MPS) [8], along with a set of techniques for organizing the manipulation of these surfaces. These two sets of tools and techniques represent two different perspectives on the user interface.

In the former case, we refer to the surfaces as *simple presentation surfaces*, or SPS's. SPS's control the user's display and user interaction with that display.

In the latter case, a *control-handler*, or *Extended Presentation Surface* (EPS), organizes and controls the activation of an SPS, as well as other EPS's. Since the control knowledge implicit in an EPS is changeable and dependent on contexts arising within both the historical context of user interaction and the execution context of the application, it constitutes a yet more sophisticated kind of surface.

The user experiences each extended presentation surface as a "menu" or "form" on the terminal screen. While the SPS may be thought of as the actual menu or form, along with the reporting of the user's interactions, the EPS interprets the meaning of those user interactions to ASTA's rule base.

We will discuss the SPS and EPS later in more detail.

4.2 SIMPLE PRESENTATION SURFACES

A Simple Presentation Surface (SPS) is primarily composed of a partition of a database of record structures, (hereafter called "objects"), and a set of relations between various object attributes and their screen display representations. These representations might be textual or graphical. The graphical and internal representations of visual objects is tightly coupled, such that manipulating an attribute (record field) of an object will change its display representation in some way, and manipulating the object's display representation will produce a corresponding change in the object's attributes.

4.2.1 MPS icons

In the Multipurpose Presentation System (MPS), the icon is defined as a display representation which may be manipulated by the user to produce changes in some intermediate structure. For the ASTA user interface, these structures are objects in the MPS database. At the current time, all icons in ASTA are textual, and manipulating these icons involves adding or deleting character strings from them. At the present time, an SPS deals with object-icons of three different types, which are defined by how they are displayed and how they are manipulated by the user.

A *selection* is simply a text string with a location on the SPS, where a cursor or other pointer may be positioned by the user and some selecting action performed, (usually a *Select* key or mouse button).

An *input* defines a highlighted box of a specified size where the user may position a pointer and type text.

A *caption* is simply an arbitrary text string at a given location.

4.2.2 SPS components

The SPS must maintain a character or pixel bitmap state image of each icon, a current state bitmap of the parent display "window" on which the icon appears, a layout template of icon locations on that window, and a set of permissible manipulations for each object-icon type and for each aggregate of object-icons on the window.

4.2.3 Granularity of change reporting

Both the display objects and the application objects in ASTA interact by changing the interface database. One of the problems in user interface design is at what times to report these changes to the user or to ASTA's Extended Presentation Surfaces (EPS).

If the application alters the interface object database for a given SPS, the display representation is updated immediately. If the user manipulates an icon on the screen, the interface waits to notify the application until the user is finished interacting with that SPS, although what "finished" means may be rendered transparent to the user by the extended presentation surface techniques described below. Thus the application's EPS is notified of the results of interactions with the user in "chunks," with the granularity of the chunk set at the SPS. SPSs are displayed one at a time and the results of the interaction are delivered to the corresponding EPS for interpretation at intervals defined by the EPS.

The user terminates interaction with the object set of a given SPS by pressing one of the pre-defined legal buttons or keys while the pointer is positioned at one of the display objects maintained by that SPS. The SPS then returns a tuple consisting of the key pressed, the object the pointer was positioned at when the key was pressed, and a list of pairs, where each pair consists of an object changed and the change made by the user. The EPS might process these changes and immediately redisplay the SPS, if necessary, resulting in no real "termination" of interaction as far as the user is concerned.

4.3 MPS SUPPORT FOR SIMPLE PRESENTATION SURFACES

The Multipurpose Presentation System (MPS) provides support tools for a set of Simple Presentation Surfaces (SPS's):

- Display Engine
- Cursor locator and positioning device
- Object finder
- Database of objects and object representations
- Update relations between display representations and objects
- Query and update by application control structures (EPS's)
- SPS activator and interpreter

We describe each of these below in more detail.

4.3.1 Display Engine

An icon's display representation can be placed on a screen at a given location. MPS also supports a window manager built on top of the Berkeley Software Division's CURSES package. This display engine, while limited to textual windows, provides many capabilities similar to graphics window managers. The CURSES package allows the implementation of the ASTA user interface on a wide variety of terminal types.

4.3.2 Cursor locator and positioning device

CURSES provides facilities for moving and locating the cursor within different windows. MPS maps keyboard control characters, which may be programmed into function keys, onto these functions.

4.3.3 Object finder

Given a cursor or pointer location, MPS will return an object from its database which has a location corresponding to the position. Since MPS currently restricts the locations to which the cursor can move to those locations within textual icon objects, this lookup is greatly simplified. This has proven an advantage on the heavily loaded time-sharing machines where ASTA was developed.

4.3.4 Database of objects and object representations

MPS uses the Franz Lisp "defstruct" structure as a basis for its database of objects. The defstruct functionality allows the interface to define a generic record structure along with macros to create and modify instances of that structure. Defstruct is used in the implementation of object-icon types and at the SPS definition level. Each icon-object contains information about its parent surface, location, current display representation appearance and current "value." MPS maintains knowledge on how to draw the display for those icon-object types that it supports.

4.3.5 Update relations between display representations and objects

When a display representation is altered by the user, MPS updates the appropriate data base object. Similarly, when a data base object is changed by a call from an EPS (see below), the display representation bitmap is revised according to the object's type.

4.3.6 Query and update by application control structures (EPS)

MPS allows external query and update of its database of objects. For the ASTA user interface, the extended presentation surfaces are the agencies that make these calls, which can be made at any time. It is possible for a given icon-

object to be a member of several different SPS's.

4.3.7 SPS activator and interpreter

MPS maintains a central SPS activator to display the SPS window and handle interaction with the SPS's icon-objects. It is an interpreter in the sense that it "parses" the icon manipulations done by the user. A tuple of changes made by the user is returned when the user leaves (deactivates) the current SPS for another. Each tuple contains the key pressed which caused the SPS activator to return, the icon that the cursor was positioned at when the key was pressed and a list of changes made by the user to the icons in the SPS since the SPS was activated. The SPS activator is the primary means of invocation of an SPS by an EPS.

4.4 CONSTRAINTS ON THE INTERFACE IMPOSED BY THE CURRENT IMPLEMENTATION

As is the case with any initial implementation of a research design, our ASTA prototype was subject to some implementation-dependent limitations. These are described in the following subsections.

4.4.1 Display Engine

MPS uses the Berkeley Unix 4.2 CURSES screen updating and cursor movement optimization package. CURSES uses the Berkeley UNIX 4.1 and 4.2 "termcap" library of terminal types. While this approach allows us a great deal of flexibility in implementing ASTA on a large variety of terminals, it restricts ASTA to character devices. To implement graphics, we would need to abandon CURSES and move to a graphics package that used a Graphcap-like library of graphics devices if we wished to maintain device-independent capabilities.

4.4.2 Inter-Process Communication

The ASTA interface makes use of MPS in a two-process loosely-coupled architecture. This offers the advantage of decoupling the two principal computational tasks, which makes a distributed implementation easy to achieve. For the foreseeable future, MPS will make use of the system facilities to support "pipes" under UNIX 4.1bsd and 4.2bsd, however, which means that some reimplementation will be necessary to support the message delivery mechanisms currently in use by ASTA before it can be installed on a two-node distributed system.

4.4.3 Defstruct data structures

The MPS database makes extensive use of the "defstruct" record structure in Franz Lisp. The support for this structure is complex and available in the following dialects of Lisp: Franz Lisp, ZetaLisp and Common Lisp. Current plans are to make even greater use of LISP structures and ZetaLisp style "flavors," while retaining Franz.

4.4.4 C Language Interface

Franz Lisp is one of the few dialects of Lisp that allows direct access to "C" language functions. ASTA, as well as the MPS interface, makes extensive calls to C functions in external packages such as CURSES. In addition, while the great majority of ASTA itself is written in LISP, some C code was necessary to overcome design and implementation flaws in Franz. At the time of this writing, Franz Lisp runs only under the Berkeley versions of UNIX.

4.5 ASTA--MPS ARCHITECTURE OVERVIEW

ASTA makes use of the MPS user interface in a two-process architecture under the UNIX operating system. Communications between these two processes are implemented under the UNIX 4.1bsd version of "pipes."

MPS incorporates the total set of SPS's in pre-compiled form, since these seldom change in the course of a working session. The MPS system and these pre-compiled SPS's are kept in an executable Franz Lisp "dumplisp" image, called "asta."

At run time, the user starts up the MPS image, which in turn executes the process containing FranzLisp, MRS, the rule base, the EPS interpreter, and other ASTA code, called "CMRS" (for "Control-MRS"). A simple first EPS is activated from which the user is instructed to choose a subsystem, and the interaction begins.

Under normal conditions, interprocess communication consists of the EPS-ASTA process sending remote function calls to the MPS process, and the MPS process returning values. Calls may be synchronous or asynchronous as needed. The user never talks to the CMRS process directly, but talks to the asta process.

The ASTA designer has several simple tools to aid in the design and debugging of new subsystems:

- *Debugger*: MPS allows the Franz Lisp language debugger to be invoked remotely from the MPS process if an error condition arises.
- *Logs*: Both the ASTA and the CMRS processes maintain running logs of their interprocess communication and error states in text files in the ASTA directory. If an error state is detected, as complete a description

of the error as is possible will be written into the appropriate log file.

- *Direct access:* The designer may bypass the EPS interface and invoke ASTA functions without going through the usual SPS directly, through a special subsystem interface. This permits the designer to distinguish between behavioral characteristics (including potential errors) of the user interface and the inference engine, aiding the fault isolation and program development process.

4.6 EXTENDED PRESENTATION SURFACES (EPS)

An Extended Presentation Surface (EPS) is responsible for maintaining the relationships between one or more Simple Presentation Surfaces (SPS) and the corresponding ASTA functionality.

The ASTA user interface maintains several EPS interpreters, which activate the corresponding SPS interpreter and handle responses. An EPS is in turn activated by one of 12 form and menu "handlers." These handlers in turn activate the EPS interpreters in the new context of a different form or menu, or make calls to the appropriate MRS functionality.

When a handler returns, control passes to the EPS that called it. When an EPS returns, control passes to the handler that called it. All handlers and EPS's run in the context of a top-level control handler, which never returns until the program is terminated by the user.

The EPS interpreter makes use of the programmable function keys on the Ann Arbor Ambassador display terminal.

4.7 REPORT GENERATION FACILITY

From any form, the user may press the REPORT key, and all of the current data in all the forms in the system will be dumped to a time-stamped file. The report generator may be invoked any number of times during a session. The format of the file name is *"day-month.hour-minute.rpt"*.

4.8 CHECKPOINTING FACILITY

The user can save the state of his analysis at any point, end the session, start another session at a later time, and take up where he left off by restoring the saved analysis.

The analysis can be saved under the default name "analysis," or under a name that the user may choose. Analyses are saved in two files, an MPS section and an MRS section. When the analysis is restored, the MPS section restores the state of the user interface, while the MRS section restores the state of the knowledge base.

The user may ask for a list of all the analyses saved to date in the local directory.

It is also possible to merge multiple analyses using the "merge" option of the checkpoint restoral mechanism. This may be desirable when, for example, an analyst has performed an analysis of the antenna subsystem of one radar system and the transmitter of another radar system, and later discovers, hypothesizes, or decides that the antenna and transmitter are in fact from the same radar system. By merging the subsystem analyses, the radar system analyst can determine what the interaction of the subsystems is likely to be and draw additional inferences about the capabilities of these and other subsystems of the radar system.

4.9 FUTURE DIRECTIONS FOR THE USER INTERFACE

4.9.1 Consolidation of EPS interpreters

The various EPS interpreters should be consolidated to one interpreter for forms and menus. This would considerably simplify the design and implementation of new ASTA functionality. We would incorporate a small database of context information into each handler, so that the handler would know something about where it fit into the control hierarchy. The top level loop would then consist of repeated calls to the same EPS interpreter. Each handler would be responsible for maintaining its own context information. Under this scheme, there would be no control nesting of the handler or interpreter calls, and the user could move directly from, say, signal specification to examination of the radar transmitter form without climbing back up, and then down, the menu hierarchy.

4.9.2 Subsystems

With the consolidation mentioned above, it would now be possible to partition groups of forms and menus by context. The partitions might include signal properties, physical components, help, system utilities, and others.

There is no limit on the number of subsystems that ASTA may incorporate. Every effort has been made to keep these simple and avoid unnecessary subordinate subsystems. While the subsystem structure could be easily modified to accomodate subordinate subsystems to any depth, we believe that the user should not be burdened with such complexity at the subsystem level, and that the number of subsystems should be limited.

4.9.3 Help

Help should be available at several levels. Help might be required about the user's task at hand (*static documentation*), where they are in a subsystem (*orientation help*), what their options are (*static contextual help*), or what they should do given either their interaction with the system thus far or their stated goals (*dynamic contextual help*). Static documentation can be compiled into the SPS as static text icons or provided as a separate co-display. Static contextual help is user independent state information that can be gleaned from ASTA or found in the EPS structure. Dynamic contextual help depends on constructing and maintaining knowledge about the user's intentions. The first two kinds of help could be implemented through the HELP subsystem mentioned above. The third type of help is a research problem in its own right.

4.9.4 SPS

ASTA's SPS's are rather dense, which reflects the self-imposed restriction that we be able to run the system on 24-line terminals. We plan to redesign some of these SPS's to allow for a larger variety of help and notification windows, as well as to allow ASTA to run on terminals with larger screens.

4.9.5 Checkpointing

The checkpointing facility currently allows the user the ability to merge several checkpointed analyses. The semantics of this operation is the simplest possible, however: one analysis is brought in "on top of" another. It is assumed that there are not conflicts between the two analyses. This is reasonable if, for example, one analysis concerns an antenna subsystem and another concerns basic signal structure. Such independence of analyses is not verified, however, before the two analyses are merged. Considerable research remains to be performed to determine what minimum conditions must be met before two arbitrary

syntactically compatible knowledge bases can be merged into one larger semantically meaningful knowledge base.

4.9.6 Constraint checking

Currently, constraint checking of user input is located in the EPS. Recent MPS development has resulted in an ability to check hard constraints (those constraints which do not change over the course of a working session) within the SPS. This will speed up the interaction, as well as affecting the static documentation aspect of help.

4.9.7 Pop-up and pull-down menus

MPS has the capability to use pop-up and pull-down menus. More effort is needed to determine how these features might be best exploited in ASTA.

4.9.8 Presentation Graphics Support

MPS is currently without graphics support. If an appropriate graphics package were found, it could replace the CURSES package currently in use, though a decision might have to be made to abandon terminal independence. If such steps were taken, it would also be desirable to acquire or write a set of graph drawing routines.

4.9.9 Mouse support

We might wish to allow mouse support to replace the keypad interaction. Again, this is a matter of replacing a module in MPS by writing or acquiring a mouse interaction package, and it might require that we sacrifice some device independence in order to gain the increased functionality.

5. CONFLICT DETECTION AND RESOLUTION

Artificial Intelligence problems are typically characterized by the imprecise, inaccurate, and sometimes incorrect data from the environment with which they must contend. Problems of this type will frequently lead to inconsistent, conflicting conclusions.

The analysis of radar signal parameters performed by the ASTA system is one example of a problem area within radar analysis that has the potential for inconsistent and conflicting results. Radar parameters supplied to the system are collected from the electromagnetic environment in a way that leads to incomplete and inaccurate measurements. Both numeric and symbolic values may produce a conflict within the system in a variety of ways. In general, this description of attempting to achieve a plausible explanation in the face of ambiguous knowledge provides an excellent characterization of the radar system reverse-engineering analysis process.

The user, however, is typically interested in "what the answer is," irrespective of the uncertainty inherent in the problem domain or data. Humans make difficult choices when faced with ambiguous or conflicting data and knowledge, and they presumably expect any other "expert" -- including an automated one -- to do the same. An approach to handling this uncertainty therefore must be employed in order to present the user with consistent solution to the reasoning problem that the radar analysis expert system is trying to solve. In this chapter we discuss mechanisms we have studied to identify and resolve such data and reasoning conflicts.

5.1 SOURCES OF CONFLICT

There are many ways in which conflict can surface within the ASTA system. They can be characterized in two classes. *Shallow conflict* involves two or more values derived for the same radar system attribute. *Deep conflict* is detected when two separate values are inconsistent in terms of the relative quantity or quality of the relationship.

Shallow conflicts are produced when there is more than one value for a single attribute. The attribute may be either a numeric or symbolic value. One source for values in the ASTA system is for the analyst to assert what he claims is true based directly on a measurement value or based on an estimate from his own incomplete data. A second source of shallow conflicts is the inference process, which automatically asserts values based on previous values entered by the analyst. For this case, the reasoning system that derived the values must then be interrogated in order to determine the source of the conflict. The rules can be traced back to the primitive values input by the analyst if the value was derived by rule applications, or the analyst can be queried directly if the source was an analyst input value. In either case, the cause of the conflict is due to an inconsistency in the primitive data gathered as input to the system under the assumption that the rules used to derive the result are correct. The implications of relaxing this assumption about rules is described in a following section. Once this inconsistency is recognized by the system, a method for resolution is employed.

Deep conflicts are a result of an inconsistency in the relationship between two values. Typically, the attribute values are related in a general way rather than by specific values. Obviously, if two attributes had a directly computable relationship, a domain rule would be used to derive this value directly. For conflict detection, the values are usually related by a qualitative measure of their relationship, such as "greater than" or "less than." For discrete values, the relationship may narrow the set of possible value assignments, rather than specifying only one possible such assignment. Examples of these types of conflicts are

appear in a subsequent section.

5.2 WHAT DEFINES A CONFLICT

Due to the imprecision of the data, the conflict detection algorithm cannot reasonably expect to consistently achieve an exact match of multiply-derived attribute values. Conflicts can arise that involve numeric attribute values or symbolic attribute values. For continuous numeric values, therefore, ASTA's algorithm allows a small range of error bounding the difference between the two values compared for conflict. The bound is based on a percentage of the current system value. Typically the percentage is a constant for all the attributes of the system but may be overridden individually for certain attributes by the analyst. The new value may be a single percentage, or it can be defined using rules within the rulebase. This allows a very general way to define the error ranges based on the current context and the importance of discovering conflicts for certain distinctive radar attributes.

A second method studied to define the bound is to propagate uncertainty bounds through the rules to derive a range on the inferred value based on the uncertainty of the values used to define it. This method is intuitively elegant but is difficult to implement with respect to the methods of combining together ranges of values within the rules to produce accurate ranges for the new inferred value. These ranges may also tend to expand to unreasonable bounds based on worst case estimates for all bounds combined which is an unlikely case.

In the case of discrete value comparison when more than one method is used to derive the attribute, only a subset of a known finite set of instances is possible. This reduces the conflict detection problem in certain respects. Generally, if a second discrete attribute value is derived in some way it is compared with the existing value and if different, a conflict is reported.

5.3 RESOLUTION OF CONFLICTS

Once a conflict has been detected, a procedure for resolution must be used so further inferences can be drawn based upon one consistent value for the attribute (that is, based upon one consistent view of the radar system under analysis). In the case where the attribute has one of its values defined directly by the analyst and the remainder derived through rule applications, the analyst defined value will be used. We employ this heuristic because the value provided by the analyst is most likely a value based directly on the data gathered for the analysis, whereas a derived value is based on analyst-entered data which have been propagated through rules that add uncertainty by virtue of assumptions and accumulated error present in the rules themselves. Where a long reasoning chain of rules has been used to derive a value, these sources of error become more pronounced, thus multiplying the uncertainty present in the derived value.

For cases wherein all the attained values are derived through rule applications, the most recently derived value is used. However, it may be more appropriate, using the logic from the previous case, to retain the value that required (say) the least number of rules to derive its value, thus minimizing the accrued error, although this is not currently implemented. A second consideration is the number of primitive input values used in the antecedents of the rules with perhaps a measure of their uncertainty computed as well. These considerations deserve thorough investigation and experimentation before determining the optimum method. This may involve a dynamic xt-dependent mechanism for choosing the method as each conflict arises. For the third case, in which the analyst has entered two values that conflict, the ASTA "change hypothesis" mechanism will remove the original attribute value and enter the new value.

The user interface alerts the analyst to the conflict; the explanation facility can then be used to provide a list of the conflicting values of the attribute.

During conflict explanation, the source of the conflicting value, which may be either analyst-provided or derived through rule application, is traced. For the derived values, a trace mechanism allows the analyst to follow the rule applications back to the input primitives used. By providing this information, the analyst may be able to determine the reason for the conflict and change the erroneous or conflicting data using the change hypothesis mechanism. This will update not only the presently conflicting value but all of the values derived with the erroneous data that may cause future conflicts. An additional benefit of allowing a trace of the conflicting value derivations is that the analyst may decide to override the retained value based on his own intuition and his understanding of the input radar data, once a more informed basis for judgement has been obtained through inspection of the system's reasoning process that led to the conflict.

5.4 IMPLEMENTATION

As implemented, conflict detection consists of four major tasks or components; the detection of both numeric and discrete multiple value assertions for a single radar attribute, detection of "deep" conflicts between related values, a conflict resolution algorithm, and interface tools to alert the analyst to conflicts and allow him to interrogate the sources of conflict.

The detection of multiple conflicting values for one attribute is performed whenever a new value is asserted into the MRS database. A procedural attachment is added to the assertion mechanism using **toassert** in conjunction with the **mps-stash** function. Ideally, these functions should be kept separate, but given the MRS constraint which only allows one **toassert** procedure to be used at any time, the two functions were combined and executed in a cooperative manner. Specifically, before a **stash** of the new value can be executed, a check for conflicts must be completed and only if none exist can the value be stashed in the MRS database and displayed on the MPS interface. The first step when

detecting multiple conflicts involves checking for the existence a value already in the database using the **lookup** MRS function. If the value does not exist, then there is no conflict; if one does exist, a comparison must be performed. The second step requires determining the type of the value: if it is a discrete type, then the new value and the old value are compared absolutely for conflicts, whereas if the type is a continuous numeric value, the values are compared within some range. The comparison method for numeric values involves a **truep** (backward-chaining) operation for the conflict range of the specific value asserted. The **truep** returns a value for the percentage error allowable between the new value and the old value. This percentage may be determined by the application of a complex set of rules or it may be a single assertion in the database. In either case, the analyst and system designers now have a flexible mechanism for determining error ranges for attributes and may be modified independently of the domain rules themselves.

The detection of "deep" conflicts in the ASTA system is implemented via the MRS forward chaining mechanism, using metarules designed specifically for detecting conflicts. When a new value is determined, a special conflict detection MRS theory containing the conflict detection rules is activated. By activating the theory, these rules now become applicable and can be run. The *sta-rules* theory, containing many of the domain rules used in the ASTA system, is deactivated so inadvertent rules aren't run when a fact is asserted into the conflict detection theory. After the appropriate MRS theory activations are completed, the new fact is asserted into a temporary data space made specifically for storing conflict detection assertions. When the conflict detection check is completed, this theory is emptied of all its data and deactivated until the next conflict check is made. When the fact is asserted it is forward chained within the conflict detection theory. If the fact matches a portion of rule, the remaining antecedents of the rule are checked for matches, and if the rule is run it may assert a conflict between two related attributes in the system. Otherwise, there are no conflict detection rules which notice any conflicts between related values.

The conflict resolution algorithm determines the appropriate value to use within the system and its implementation is tightly integrated with the MPS display and the MRS database stashing mechanisms. The conflict detection function returns a flag describing the result of the conflict detection check. If there is a conflict and the new value should be displayed and used for further inference, then the MPS display is updated and the conflict fact is stashed in the MRS database; otherwise the value itself conflicts and is determined to be the most likely value to be in error, so no action is taken other than a stash of the conflict assertion.

When a conflict is asserted into the MRS database, the user must be alerted. A conflict is represented in MRS as the 3-tuple:

(conflicting *old-assert new-assert*)

A procedural attachment for the display function is added using the *costash* mechanism for any fact stashed that matches this template. When the template is matched, the procedural attachment puts a "?" next to both the old-assert attribute and the new-assert attribute on the MPS forms. The system recognition of the conflict is therefore presented to the analyst so she may investigate the nature of the conflict and take appropriate action.

5.5 EXAMPLES OF THE USE OF CONFLICT DETECTION

In this section five typical applications of conflict detection are discussed. These applications are:

- Reasoning from multiple points of view.

- Detecting inconsistencies between derived numerical results and symbolic, qualitative assumptions.
- Detecting rule inconsistency.
- Detecting measurement errors or errors in hypothetical assertions.
- Indication of new technology.

5.5.1 Reasoning From Multiple Points of View

In many cases a particular slot on a form may be derived from alternate points of view. For example, if it is desired to estimate the maximum detection range of a pulsed doppler radar, at least three independent approaches are possible:

- Harmonic analysis of the pulse repetition frequencies may be used as described in Section 3.4 to determine the unambiguous range of a dwell. The largest unambiguous range obtained will normally be associated with the autonomous search mode and should be equal to or greater than the maximum detection range of the radar.
- The radar range equation can be used to estimate the maximum detection range, if sufficient data are available to estimate or hypothesize enough of its terms. For example, if the gain is estimated from ELINT frequency-scaled analysis of PHOTINT, or beamwidth measurements are available and radiated power has been measured, then the detection range can be calculated based on reasonable assumptions about system losses and

detection characteristics.

- If the associated weapons system can be hypothesized, then engagement analysis, as described in Section 3.4, can be used to estimate the maximum detection range, based on estimates of the interceptor missile capabilities and target characteristics.

The consistency of the results derived from these separate viewpoints improves the confidence in the underlying assumptions supporting each calculation. If two out of three results yield consistent interpretations, while a third approach yields a discrepancy, then it is likely that some of the assumptions supporting the third approach are in error.

5.5.2 Inconsistency Between Qualitative Abstractions and Numerical Calculations

Symbolic descriptions of the radar, summarizing its functional objectives or overall performance, can be compared for reasonableness with the results of numerical calculations based on measured data. For example, large tracking errors for a dwell in any dimension are not consistent with that dwell being used in the track operational mode. When such consistency checks are performed regularly, it is possible to make broad assumptions initially and withdraw them only if a conflict with the observations is found. This technique is useful for getting started on a limited data set that does not support a rigorous instantiation of any particular hypothesis, but which gives reasonable support to a general context for the analysis.

5.5.3 Rule Inconsistency

Inconsistent results may also indicate that the knowledge base contains an error. The most common error of this type is over-generalization of a rule, that is, applying it to a broader class of contexts than is warranted due to insufficient antecedents in a rule. For example, if one attempts to apply the unambiguous range to estimate the maximum detection range of a low PRF, pulsed, MTI, radar system, then the maximum detection range obtained may be much larger than the actual unambiguous range of the radar, because the purpose of pulse stagger in this case is primarily to eliminate blind speeds.

It is also possible to detect syntactical errors indirectly using the conflict detection mechanism, since the rules leading to contradictory results can be identified as candidates for correction. An error in the specification of the maximum power in a rule describing magnetron characteristics was found in this way by noting that the rule produced contrary evidence for the use of magnetrons, even though it had received positive support from several other lines of reasoning.

5.5.4 Measurement or Assumption in Error

The consistency maintenance mechanism guarantees that any errors introduced during the initial phases of the analysis will eventually be detected as data is acquired and associated with the knowledge describing the system. For example, if an incorrect frequency is used for antenna scaling due to reporting of an harmonic of the true radiated frequency, then this will conflict with other inferred frequency dependent aspects of system behavior derived from other data, indicating the possibility of error in the original measurement. The capability to detect conflicts in this way also permits default values to be used in the early stages of analysis and later withdrawn when the accumulated evidence indicates that some other value is more appropriate in the current analysis context. The incorrect association of a weapons system with the radar system would be

indicated by an inconsistency in weapons dependent analyses, such as the estimation of maximum detection range via engagement analysis on the one hand, and by evaluation of the radar range equation on the other.

5.5.5 Indications of New Technology

The rule base can be tuned to any particular set of assumptions about the state of the art of a particular technology by defining the parameters and the scope of application of certain rules in terms of postulated capabilities. One can, for example, estimate the system configuration based on different beliefs about the power capabilities of magnetrons or the level of signal processing sophistication presumed to be available. If ASTA detects a conflict in the knowledge base, it may be possible to eliminate the conflict by relaxing one or more of the technology constraints. If this is possible, it may indicate that the technology level represented by the original rule has been underestimated or has developed substantially since the rule base was last updated.

6. SUBCLUTTER VISIBILITY MODELING

In this chapter we discuss ASTA's capability for performing subclutter visibility modeling. Although a complete description of the Moving Target Indicator (MTI) subsystem was beyond the scope of this project, an approximate description of the relative effectiveness of various MTI techniques to suppress clutter was desired. The approach taken here is one currently in practical use by expert radar analysts for performing subclutter visibility evaluation.

1. The performance of the MTI system is expressed in terms of the subclutter visibility. Subclutter visibility is the ratio by which the target echo power can be less than the input clutter power and still produce a detectable output with stated probability. Thus, for example, a subclutter visibility of 20 dB implies that the radar can detect a target in clutter that has a return 20 dB (that is, 100 times) stronger.
2. Subclutter visibility is evaluated as a fourth order polynomial in the logarithm of the ratio of the standard deviation of the clutter spectrum to the pulse repetition frequency. The coefficients were supplied by experts who perform subclutter visibility analysis; these coefficients depend upon the type of MTI and the number of cancellers used.
3. Heuristics are used to estimate the type of MTI technology. MTI technology is a symbolic attribute which can assume the values *old storage tube*, *new storage tube*, *delay line*, or *digital*.
4. The number of cancellers can be equal to 1 or 2 and is input by the analyst.

5. A zero mean, Gaussian frequency spectrum is assumed for the clutter. Thus, it is assumed that any platform motion has been compensated.
6. The standard deviation of the clutter spectrum is small compared with the PRF.
7. The total clutter spread is due to a number of independent sources. The contribution of each of these sources is modeled by a Gaussian distribution. Thus the total clutter variance can be calculated by adding the variances of the individual contributors.
8. Clutter spread arises from many sources that cause frequency, phase, amplitude and envelope delay fluctuations. Three of the sources giving rise to clutter spread are modeled explicitly. The effect of the unmodeled sources is represented by specifying an upper bound on the subclutter visibility as a function of the transmitter configuration and the type of MTI. The subclutter visibility is taken as the minimum of the calculated subclutter visibility and the tabulated upper bounds.

Graphs of the subclutter visibility versus the ratio of clutter spectral standard deviation to the pulse repetition frequency are shown in Figure 6-1. The coefficients of the polynomials used to calculate the subclutter visibility are given in Table 6-1 for the single canceller case and Table 6-2 for the double canceller case. The limits on subclutter visibility imposed by the MTI-technology are shown in Table 6-3. The subclutter visibility of a modulated final amplifier transmitter configuration cannot exceed 30 dB.

The three factors giving rise to clutter spectral spreading that are explicitly modeled in ASTA are:

1. Transmitter instabilities
2. Antenna scanning, and
3. Internal motion of the scatterers.

For the intelligence analyst, a crude estimate of the probable limitations due to transmitter instabilities is sufficient in most cases. The standard deviation of the contribution to the clutter spectrum spread due to transmitter instabilities is assumed to be proportional to radio frequency. The constant of proportionality is taken as 0.0016 for master oscillator/power amplifier transmitter configurations and as 0.002933 for modulated final amplifier configurations.

The clutter spectrum broadens due to angular rotation of the antenna because the radar does not receive echoes from the identical patch of scatterers from pulse to pulse. For a Gaussian two-way antenna pattern, the clutter spectrum broadening due to antenna motion is proportional to the ratio of scan rate to antenna azimuthal beamwidth. If the beam is scanning in two directions the variances for each scan direction are added. If the beam is held stationary at each discrete beam position during the transmission and reception of an MTI pulse train, there is no contribution to the clutter spread from antenna motion.

The standard deviation of the contribution due to internal motion of the scatterers is assumed to be proportional to the ratio of the standard deviation of the relative velocity of the clutter scatterers to the wavelength. A detailed consideration of the velocity spectrum of various scatterers was not essential to the demonstration prototype. ASTA currently reports the subclutter visibility under three different assumptions about the clutter spectrum of the scatterers--ground

clutter, chaff and rain. The values used for the standard deviation of relative clutter velocity in meters/second are shown in Table 6-4. The value used for ground clutter is consistent with the clutter obtained from wooded hills with a wind speed of 25 knots. The value for sea clutter with a wind speed of 10 knots is roughly equal to the value given for chaff.

The design/development date is used to partition the MTI heuristics into two theories. For radars designed/developed prior to 1975, the following heuristics are used:

- Delay-line or storage tube MTI is assumed.
- If many stagger modes are observed, then the storage tube is probable.

For radars designed/developed after 1975, the following heuristics are used:

- Delay-line or digital MTI is assumed.
- If many stagger modes are observed, then digital MTI is more likely.
- If all PRIs and RFs are harmonics/sub-harmonics of a common clock frequency that is consistent with computer processors, then digital MTI is more likely.

The 1975 date for digital MTI is questionable. A firm date for digital MTIs is

probably after 1980. Hybrid digital and delay line MTIs are also possible.

Figure 6-2 illustrates the results of applying the subclutter visibility analysis module to a pulsed oscillator (MFA) radar with new storage tube, double cancellation MTI technology. A number of the slots on the form are filled in by other ASTA analysis modules. For example, the azimuthal beamwidth is obtained from the antenna analysis form. PRF, RF, wavelength and scan are obtained from information entered on the frame, dwell and sequence forms. The number of hits per scan is obtained by combining information from waveform and antenna analysis. Design date, transmitter configuration and MTI technology were entered by the analyst. If the analyst had entered information that was inconsistent with information that had already been computed by the transmitter analysis module or the MTI technology heuristics, then the conflict detection mechanism would have notified the analyst of the inconsistency. The computed subclutter visibility for this radar in ground clutter, chaff and rain are shown in Figure 6-1.

From the results shown in Figure 6-2, the analyst can determine that, for ground clutter, the unmodeled system instabilities are the dominant factor, since the calculated subclutter visibility reaches the heuristic upper bound. This may suggest to the analyst that a more detailed consideration of the hardware is required. Similarly, the analyst can determine that internal scatterer motion dominates the subclutter visibility for both chaff and rain. In rain the subclutter visibility will be reduced to 4dB.

The sensitivity of subclutter visibility to assumptions about MTI technology and transmitter configuration can be investigated by using the "Change Hypothesis" mechanism as discussed in Chapter Five.

Table 6-1: Single canceller subclutter visibility coefficients

Coefficient	Old Storage Tube	New Storage Tube	Delay-Line or Digital
a_0	6.50	2.62	-12.10
a_1	30.90	20.50	-13.90
a_2	33.60	28.10	3.51
a_3	9.68	8.51	0.87
a_4	1.03	0.94	0.079

AD-A160 552


SCIENCE AND TECHNOLOGY ANALYST'S ASSISTANT REVISION(U)
ADVANCED DECISION SYSTEMS MOUNTAIN VIEW CA
A S CROMATY ET AL. JAN 86 TR-1035-01 N00014-82-C-0602

2/2

UNCLASSIFIED

F/B 9/2

NL






Table 6-2: Double canceller subclutter visibility coefficients

Coefficient	Old Storage Tube	New Storage Tube	Delay-Line or Digital
a_0	9.25	11.40	-9.12
a_1	52.00	62.00	-13.30
a_2	52.70	70.40	37.60
a_3	13.10	21.20	11.50
a_4	1.19	2.34	1.29

Table 6-3: Technological limits on subclutter visibility

MTI technology	Maximum Subclutter Visibility (dB)
Old Storage Tube	20
New Storage Tube	25
Delay Line	30
Digital MTI	40

Table 6-4: Default values for clutter velocity standard deviation (m/s)

Clutter type	Velocity Standard Deviation (m/s)
ground	0.12
chaff	1.0
rain	3.24




Figure 6-1: Subclutter visibility vs. MTI, cancellers, and transmitter type

Design Date	1973
Azimuthal Beamwidth (degrees)	* 2.40
Pulse Repetition Frequency (Hz)	* 730.
Radio Frequency (MHz)	* 2620.
Wavelength (meters)	* 0.11
Scan Rate (revolutions per minute)	* 4.0
Hits Per Scan	* 73
Transmitter Configuration	mfa
MTI Technology	new storage tube
number of cancellers	2
Clutter Spectral Spread (Hz)	
Transmitter Instability	* 7.68
Antenna Scanning	* 2.67
Internal Motion	
Ground Clutter	* 2.10
Chaff	* 17.47
Rain	* 56.59
Total	
Ground Clutter	* 8.40
Chaff	* 19.27
Rain	* 57.17
Clutter Attenuation (dB)	
Ground Clutter	* 34.40
Chaff	* 20.09
Rain	* 3.77
Maximum Subclutter Visibility (dB)	* 25.00
Subclutter Visibility	
Ground Clutter	* 25.00
Chaff	* 20.09
Rain	* 3.77

Figure 6-2: Subclutter visibility example

7. EXAMPLE PROCESSING SCENARIOS

This chapter describes two typical analysis scenarios in order to familiarize the reader with some examples of how ASTA can be used. The first scenario describes the analysis of a pulsed doppler, phased-array, fire control radar. In this scenario, the incremental nature of the analysis is emphasized by showing the state of analysis of the system after a preliminary ELINT report is received, in which only the search component of the waveform is observed, following a PHO-TINT report yielding antenna measurements, and finally after a second ELINT report is received in which the system is observed multiplexing the original search waveform with a new track waveform.

The second scenario involves the analysis of a pulsed, mechanically rotating, early warning radar. The radar employs PRI stagger to eliminate the blind speeds that are a consequence of its low PRF. The radar is assumed to employ non-coherent MTI in order to decrease its similarity with the coherent, pulsed doppler radar of the first scenario. Thus the capability of ASTA to analyze radars employing state-of-the-art technology, such as phased arrays and pulsed doppler signal processing, as well as to analyze radars based on pre-1970 technology, is demonstrated. The purpose of including the second example is to demonstrate the wide range of radar systems that can be analyzed by ASTA, to illustrate how ASTA can be used to illustrate the analogies between such concepts as discrete-multiple PRF and PRI-stagger, to describe how the scan characteristics of the search antenna induce a frame-dwell-sequence structure on the waveform that is analogous to that of the electronically scanned antenna and to show the flexibility of ASTA in processing the available data in an opportunistic fashion.

In both cases, the actual outputs produced by ASTA are shown at various stages of the processing. Explanations of the salient results are contained in the text. It is not possible to illustrate some of the more desirable features of the system in a textual format -- form browsing, tracing of an explanation through several layers of development, direct interaction with the inference engine for rule modification and development, or the ease with which a report may be restored and the implications of a change in assumptions explored.

The radars used in the scenarios were developed by an engineer with only a modest familiarity with radar analysis and synthesis techniques using ASTA incrementally. In order to keep the discussion unclassified, they do not correspond to any known Soviet radars. However, the analysis approach illustrated here should assist the intelligence analyst in getting started to use the system by providing an initial set of application techniques.

7.1 ANALYSIS SCENARIO 1: PULSED DOPPLER FIRE CONTROL RADAR

The radar chosen for analysis in scenario 1 is a multifunction (search/track) pulsed doppler fire-control radar, which employs an electronically scanned, phased array antenna. Multiple discrete PRF techniques are used to resolve range ambiguities in the search mode. Phase-shift keying with pulse compression ratio of 10:1 is used to improve range resolution in the track mode. The observations consist of waveform characteristics as determined from ELINT measurements, antenna shape and dimension information as determined from PHOTINT and some assumptions about the intended target, in this case assumed to be a cruise missile, and the maximum effective range and velocity of the interceptor associated with the radar. The analysis objectives for this scenario are:

1. Determine the beam parameters of the antenna

- gain
- beamwidth
- sidelobe level

2. Determine the transmitter configuration and characteristics

- output tube type
- bandwidth
- output power

3. Determine the performance of the radar in the search mode

- resolution
- maximum detection range

4. Determine the performance of the radar in the track mode

- target sampling rate
- tracking loop bandwidths
- measurement accuracy

5. Determine weapons system characteristics

- target handling capacity

To illustrate the incremental nature of ASTA analysis applications, the data are assumed to be collected at three different times. The first collection phase involves an initial ELINT collection on the search waveform only during the early test phases of the radar development cycle. Using only this information, the following analyses are illustrated:

- ambiguous range analysis
- transmitter analysis
- preliminary system analysis

In the second phase of processing, PHOTINT estimates of the antenna shape and size are assumed to be available in addition to the initial ELINT report. The following analyses are illustrated:

- gain beamwidth analysis
- effect of changing assumptions about aperture illumination distribution function

- use of the radar range equation to estimate transmitter output power

Finally, in the third phase of analysis a second ELINT collection on the multiplexed search and track waveforms is obtained. The additional information is used to obtain:

- range, range-rate and angular resolution and measurement accuracies
- additional system level inferences
 - target handling capacity
 - maximum detection range

Finally, ASTA's use of engagement analysis to check the consistency of the radar performance estimates with the assumptions about intended target and associated interceptor is illustrated.

The search waveform consists of a cyclic repetition of three medium PRF pulse trains of roughly equal duration, duty cycle and radio frequency, but having PRFs chosen in such a way to increase the unambiguous range and velocity. Figures 7-1 through 7-3 show the basic parameters of the three pulse trains. In these figures the analyst has entered the data from the ELINT report in all slots of the frame forms that are not marked by an asterisk. All slots that are marked by an asterisk have been obtained by ASTA's inference mechanisms. The search frames are equivalent except for the PRF and pulse width. The PRF takes on values of 15, 20 and 24 kHz and the pulse duration is adjusted to obtain a

```

Advanced Information & Decision Systems                                ASTA EXPERIMENTAL SYSTEM
-----
_Help          _Debug modes      _MRS direct    _Interceptor  _Quit ASTA
_System        _Transmitter      _Antenna       _Target       _Range equation
_Sequences     _Dwell types      _Frame types   _Checkpoint   _Match table
-----
Frame Types Menu:  _Add frame type  _Accuracy table  _Return to top menu
_f1 _f2 _f3
-----
FRAME TYPE:  f1
RF (GHz) ..... 10.0      PRF mode (high/medium/low) .....*medium
PRF (KHz) ..... 15.0      Duty cycle (%) .....*9.0
PRI (us) .....*66.66666  Unambiguous range (Km) .....*10.0
Pulse shape ..... rect   Unambiguous velocity (m/s) .....*225.0
Pulse duration (us) .... 6.0  Pulses per frame .....*75.0
Pulse train duration(ms) 5.0  Waveform energy (joules) .....
Frame duration (ms) .... 5.8  Min range tracking loop bndwidth(Hz)*0.354998
Probable oper. mode..... search  Min range tracking accuracy (m) ...
Pulse modulation ..... unmod  Max range tracking accuracy (m) ...
Frequency excursion(MHz)      Min SNR to avoid track loss (dB) ..
Chip duration (usec).....

```

Figure 7-1: Pulsed Doppler scenario: Parameters for search frame f1

Advanced Information & Decision Systems

ASTA EXPERIMENTAL SYSTEM

_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

Frame Types Menu: _Add frame type _Accuracy table _Return to top menu
_f1 _f2 _f3 _f4 _f5

FRAME TYPE: f2

RF (GHz)	10.0	PRF mode (high/medium/low)	
PRF (KHz)	20.0	Duty cycle (%)	*9.0
PRI (us)		Unambiguous range (Km)	
Pulse shape	rect	Unambiguous velocity (m/s)	
Pulse duration (us)	4.5	Pulses per frame	
Pulse train duration(ms) 5.0		Waveform energy (joules)	
Frame duration (ms)	5.8	Min range tracking loop bndwdth(Hz) *0.409916	
Probable oper. mode.....	search	Min range tracking accuracy (m) ...	
Pulse modulation	unmod	Max range tracking accuracy (m) ...	
Frequency excursion(MHz)		Min SNR to avoid track loss (dB) ..	
Chip duration (usec)....			

Figure 7-2: Pulsed Doppler scenario: Parameters for search frame f2

Advanced Information & Decision Systems

ASTA EXPERIMENTAL SYSTEM

_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

Frame Types Menu: _Add frame type _Accuracy table _Return to top menu
_f1 _f2 _f3 _f4 _f5

FRAME TYPE: f3

RF (GHz) 10.0	PRF mode (high/medium/low)
PRF (KHz) 24.0	Duty cycle (%) *9.12
PRI (us)	Unambiguous range (Km)
Pulse shape rect	Unambiguous velocity (m/s)
Pulse duration (us) 3.8	Pulses per frame
Pulse train duration(ms) 5.0	Waveform energy (joules)
Frame duration (ms) 5.8	Min range tracking loop bndwidth(Hz) *0.446077
Probable oper. mode..... search	Min range tracking accuracy (m) ...
Pulse modulation unmod	Max range tracking accuracy (m) ...
Frequency excursion(MHz)	Min SNR to avoid track loss (dB) ..
Chip duration (usec)....	

Figure 7-3: Pulsed Doppler scenario: Parameters for search frame f3

constant duty cycle of approximately 9%. The medium PRF mode for frame 1 indicates that the unambiguous range is less than the maximum detection range obtained from engagement analysis and that the unambiguous velocity is less than the assumed target maximum velocity. Frame 1 is shown after several inferences have been obtained. Frames 2 and 3 show only the raw data and the inferences obtained by forward chaining on the input data.

Each spatial dwell consists of a single cycle of PRF switching. The dwell form is shown in Figure 7-4. Note that all the properties of dwell except its basic construction have been inferred by appropriate combination of the frame attributes. The most significant conclusion on this frame is the evaluation of the unambiguous range at 150 km and the unambiguous velocity at 1800m/s. This implies that this dwell can resolve range and range-rate ambiguities for the cruise missile target at detection ranges consistent with the maximum effective range of the interceptor. This provides further supporting evidence for the conclusion that this dwell is associated with the search operational mode.

Figure 7-5 explains pictorially how the dwell, unambiguous range, and velocity are calculated. The unambiguous velocity is obtained by scaling the dwell fundamental PRF normalized by the dwell nominal frequency. The dwell fundamental PRF is obtained from the component frame PRFs by least common multiple analysis. The dwell nominal frequency is obtained from the component frame RFs by averaging the smallest and largest RFs. The unambiguous range is obtained by scaling the dwell ranging PRI, which is found by performing least common multiple analysis on the frame PRIs.

The analyst next proceeds with the transmitter subsystem analysis by displaying the appropriate form. Figure 7-6 displays the results of the analyst pressing the infer button on cross-field-amplifier, klystron, and magnetron likelihoods.

```

Advanced Information & Decision Systems                                ASTA EXPERIMENTAL SYSTEM
-----
_Help          _Debug modes      _MRS direct    _Interceptor  _Quit ASTA
_System        _Transmitter      _Antenna       _Target       _Range equation
_Sequences     _Dwell types      _Frame types   _Checkpoint   _Match table
-----
Dwell Types Menu:  _Add dwell type  _Accuracy table  _Return to top menu
_d1
-----
DWELL TYPE:  d1
Sequence of frames in this dwell
Frame-type Repetitions Phantom-PRF
1. f1          1
2. f2          1
3. f3          1
4.
5.
6.
7.
8.
9.
10.
11.
Max pulse train duration (msec)*5.0
Unambiguous range (km).....*150.0
Unambiguous velocity (m/s) ....*1800.0
Probable operational mode .....*search
Min range tracking accuracy (m)
Max range tracking accuracy (m)
Fundamental PRF (KHz).....*120.0
PRF minimum (KHz).....*15.0
Ranging PRF (KHz).....*1.0
Nominal Frequency (MHz).....*10.0

```

Figure 7-4: Pulsed Doppler scenario: Parameters for search dwell d1

AMBIGUITY RESOLUTION
DISCRETE MULTIPLE - PRF

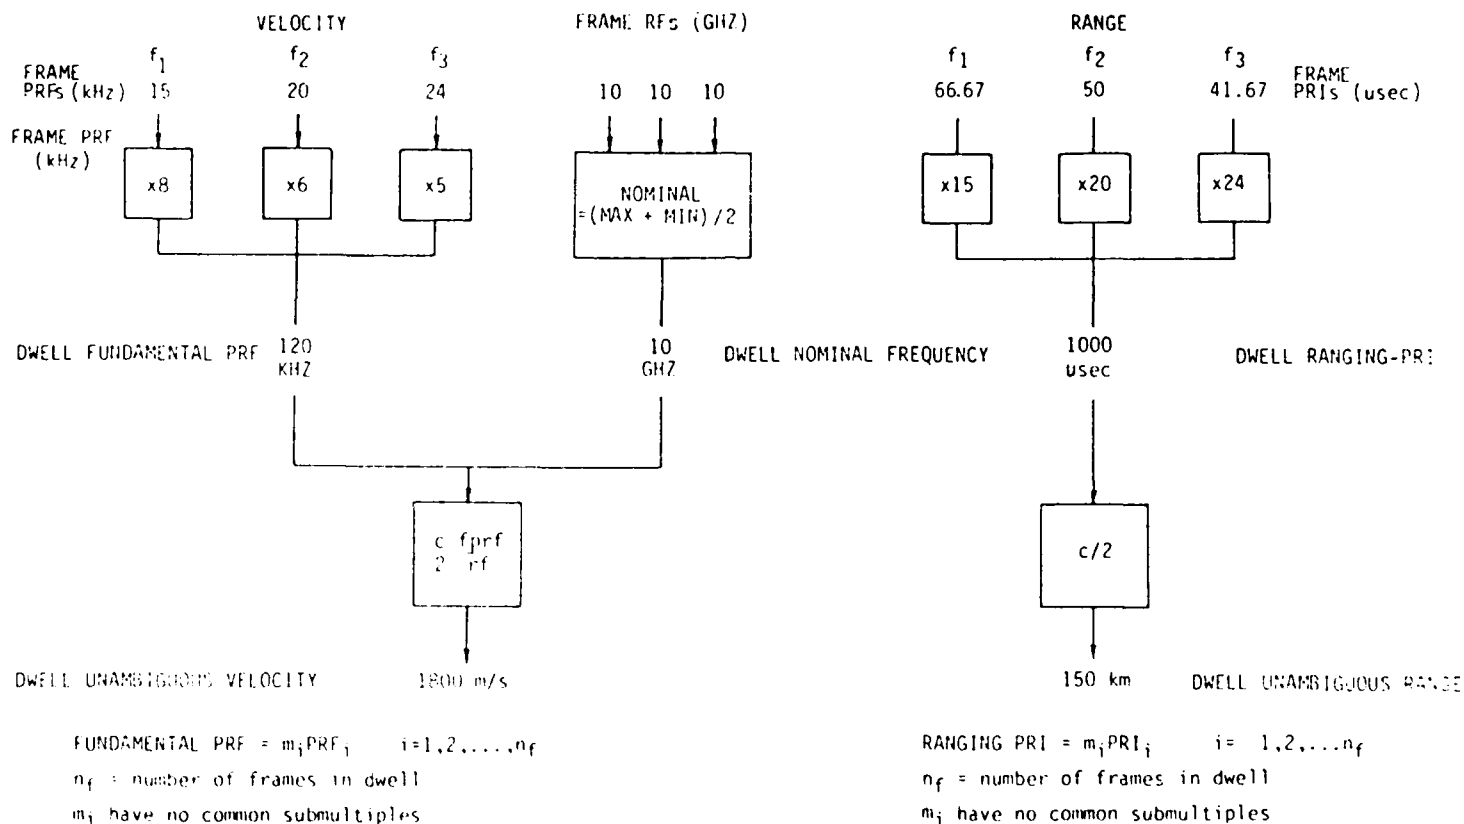


Figure 7-5: Calculation of dwell unambiguous range and velocity
for search dwell d1

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

RADAR TRANSMITTER

Oscillator configuration (mopa,mfa).....*mopa

Peak power output (KWatt)

Average power output (KWatt)

Nominal frequency (GHz)*10.0

Operating frequency band

Fractional instantaneous bandwidth.....*0.001

Coherent? (yes/no)*yes

Instantaneous bandwidth (Mhz).....*10.0

Maximum duty cycle*10.96

Minimum duty cycle*9.0

Tube type likelihoods:

TWT likelihood*

CFA likelihood*likely

Klystron likelihood*likely

Magnetron likelihood*unlikely

Figure 7-6: Pulsed Doppler scenario: Transmitter form

Table 7-1: Logical inference structure for transmitter output tube type

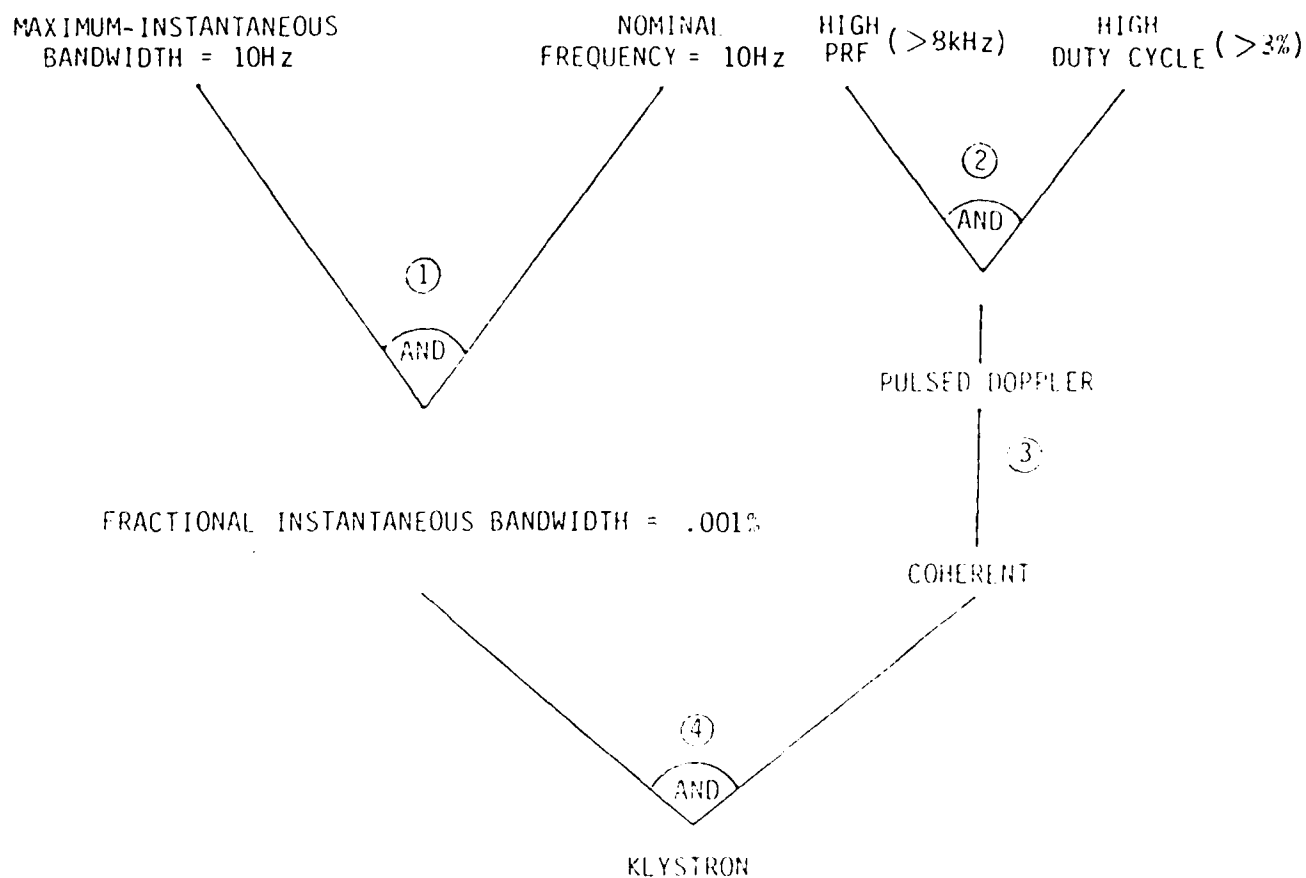


Table 7-1 illustrates how the observed waveform characteristics were combined with engineering rules of thumb and a few simple calculations to establish that a klystron would be a likely choice of transmitter output tube for this scenario. The presence of high PRF (>8 kHz) and high duty cycle ($> 3\%$) in every frame was used as evidence of pulsed doppler signal processing, which was then used to establish that a coherent source was required. The instantaneous bandwidth of each frame was calculated, using the inverse of the effective pulse duration. For unmodulated pulses the effective pulse duration is equal to the pulse duration; for PSK-modulated pulses the effective pulse duration is equal to the chip duration; for linear-fm modulated pulses the effective pulse duration is equal to the inverse of the swept bandwidth. The maximum fractional instantaneous bandwidth was obtained by taking the ratio of the maximum frame bandwidth to the nominal frequency. Finally, the transmitter designer's rule of thumb that for a narrow-bandwidth, coherent system, the klystron might well be the preferred approach, is used to reach the final conclusion. All of the intermediate results of inferences and calculations are displayed on the appropriate form.

The four primary rules used to reach the conclusion that the klystron is a likely output tube type are shown in Table 7-2. The lower level frame parameter processing rules used to establish the maximum instantaneous bandwidth and the nominal frequency are not shown for the sake of brevity. Rule 1 states that the fractional instantaneous bandwidth can be calculated from the ratio of the instantaneous bandwidth to the nominal frequency. Rule 2 is enabled whenever the instantaneous bandwidth and the nominal frequency have been previously calculated and stored in the collection of facts describing the current state of knowledge of the system. Thus the rule contains not only the expression used for calculating the result, but also the more general information that the expression can only be solved for the independent variable, if the two independent variables are known. Rule 2 states that, if the number of elements in the set of signal frames satisfying the high duty cycle and high PRF condition is equal to the total number of signal frames (thereby establishing that every signal frame satisfies the

Table 7-2: Rules that lead to klystron as output tube type

```

1 (stash-rule%explanation
  '(if (and (inst-bandwidth radar-xmtr $x)
            (nominal-frequency radar-xmtr $y)
            (is (/ (/ $x $y) 1000.0) $z))
    (frac-inst-bandwidth radar-xmtr $z))
  "The fractional instantaneous bandwidth is the instantaneous
  bandwidth divided by the nominal frequency of the transmitter.")

2 (stash-rule%explanation
  '(if (and (setof $n (and (radar-signal-frame $n)
                          (prf $n $a)
                          (> $a 8)
                          (duty-cycle $n $b)
                          (> $b 3))
                          $q)
    (setof $m (radar-signal-frame $m) $p)
    (length $p $c)
    (length $q $d)
    (= $c $d))
    (modulation-type radar-system p-dop))
  "If every frame observed has:
  a PRF greater than 8 KHz and
  a duty cycle greater than 3%,
  then the radar system employs pulsed-Doppler modulation.")

3 (stash-rule%explanation
  '(if (modulation-type radar-system p-dop)
    ((coherency radar-xmtr yes) ))
  "Pulsed-Doppler systems normally employ coherent modulation.")

4 (stash-rule%explanation
  '(if (and (frac-inst-bandwidth radar-xmtr $x)
            (< $x .05)
            (coherency radar-xmtr yes))
    (tube-type radar-xmtr klystr ))
  "The klystron is the preferred transmitter tube type
  for a narrowband, coherent system.")

```

condition), then the radar system employs pulsed doppler modulation. This example shows that the conditions for triggering a symbolic conclusion can (1) depend on an arithmetically calculated precondition and (2) the set of objects over which the constraint is satisfied can be specified in the rule. Rule 3 is of the purely symbolic type commonly employed in expert systems. Rule 4 illustrates how the satisfaction of a symbolic constraint and an arithmetic constraint can be used as preconditions for a symbolic conclusion.

The klystron likely inference example illustrates how ASTA can employ a wide variety of mixtures of symbolic and arithmetic expressions to obtain its conclusions, and also how two parallel reasoning paths (pulse characteristics to infer bandwidth and PRF to infer pulsed doppler) can be combined to produce the desired conclusion (klystron tube type).

Figure 7-7 shows the radar system summary form after processing the first ELINT collection. It is known that the system is a pulsed Doppler system with a maximum detection range of approximately 150 km. The maximum detection range estimate is obtained from the ambiguous range of the search dwell (this inference is made only for pulsed doppler systems) and the consistency of this estimate with target and interceptor characteristics in an engagement analysis. Properties associated with the track mode (accuracy, target handling capacity) can not be obtained at this time. There is insufficient data to determine the peak ERP.

The second data collection is assumed to consist of a PHOTINT observation from which antenna shape and dimensions can be approximated. Figure 7-8 illustrates how this data is entered on the form. The planar type specifies that aperture analysis is to be applied. The shape is used to determine the applicable class of illumination functions and for area calculations. The analyst makes an initial guess that the aperture distribution is uniform and presses the fill key. Using the first ELINT collection to scale the antenna dimensions in wavelengths and using the table lookup Fourier analysis scheme discussed in Section 3.3, the gain, beamwidths and first sidelobe level are calculated as shown. The analyst

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

RADAR SYSTEM

Nationality (us/ussr)

Operational modes (s,t,mixed)...

Modulation type (pulsed,p-dop)...*p-dop

Unambiguous range (km)*150.0

Unambiguous velocity (m/s)

Max range track loop bndwdth(Hz)

Min range track loop bndwdth(Hz)

Tracking accuracy

 Velocity (km/s)

 Azimuth (deg)

 Elevation (deg)

prob. target handling capacity .

Peak ERP (dBwatt)

Target revisit interval (msec) ..

Maximum detection range (km)*150.0

Figure 7-7: Pulsed Doppler scenario: Radar System summary form
after first ELINT collection

Advanced Information & Decision Systems		ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor
_System	_Transmitter	_Antenna	_Quit ASTA
_Sequences	_Dwell types	_Frame types	_Target
			_Range equation
			_Checkpoint
			_Match table
RADAR ANTENNA			
Type(planar, horn+refl).....	planar	Beam Characteristics	
Physical Description		beam width - az (deg)	*0.762
shape(rect,circ).....	rect	beam width - el (deg).....	*0.6096
Dimensions		gain (dB)	*48.4388
height (m)	2.5	first sidelobe level (dB) ..	*-13
width (m)	2.0	efficiency	*1.0
radius if horn+r (m) ...		nominal wavelength(m).....	*0.03
Area (m^2)	*5.0	polarization control(yes/no)..	*yes
aperture			
distribution.....	uniform		

Figure 7-8: Pulsed Doppler scenario: Antenna form
(uniform aperture illumination)

determines that the -13dB first sidelobe level is too large for the intended application and chooses a smoother illumination function, cosine-squared, as an alternative. ASTA removes all derived and inferred values that depend on the illumination function and replaces them with values calculated or inferred under the new assumed illumination function.

As discussed in "Antenna" section, the smoother illumination function results in beam broadening, gain/efficiency reduction and a lower first sidelobe level (see figure 7-9). Inferences on other forms that depend on the illumination function are also recalculated. For example, if the antenna gain had been used in the radar range equation, the radar range equation would be recalled with the new calculated gain, and, if angular resolution or accuracy had been calculated it would be recalculated for the new beamwidths. Inferences that do not depend on the changed data are not recalculated. For example, if the analyst had changed the width in Figure 7-8, the elevation beamwidth would not be recalculated, but the remaining beam characteristics would be recalculated.

Next, it is shown how the radar range equation form may be used to obtain an estimate of the peak and average transmitter output power. The state of the radar range equation form after pressing the fill button is shown in Figure 7-10. Five of the ten parameters in this radar range equation have been inferred from previously entered data. Three of the inferences depend only on ELINT data (maximum detection range, integration time and wavelength). The antenna gain estimate depends on both the ELINT and PHOTINT. The radar cross section was posted on the antenna form based on assumptions entered on the target form.

ASTA has also determined the form of the radar range equation to use from the pulsed doppler inference. In the pulsed doppler case integration time is related to pulse train duration and the calculated power is average power. The analyst determines that he would like to estimate the peak power of the radar by supplying estimates of the remaining four parameters. A summary of this use of the radar range equation is given in Figure 7-11.

Advanced Information & Decision Systems		ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor
_System	_Transmitter	_Antenna	_Target
_Sequences	_Dwell types	_Frame types	_Checkpoint
			_Quit ASTA
			_Range equation
			_Match table

RADAR ANTENNA		Beam Characteristics	
Type(planar, horn+refl).....	planar	beam width - az (deg)	*1.248
Physical Description		beam width - el (deg)	*0.9984
shape(rect,circ).....	rect	gain (dB)	*46.6996
Dimensions		first sidelobe level (dB) ..	*-31
height (m)	2.5	efficiency	*0.67
width (m)	2.0	nominal wavelength(m).....	*0.03
radius if horn+r (m) ...		polarization control(yes/no)..	*yes
Area (m ²)	*5.0		
aperture			
distribution.....	cos-sq		

Figure 7-9: Pulsed Doppler scenario: Antenna form
(cosine-squared aperture illumination)

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

Radar range equation parameters

Average power (Kw)
 Peak power (Kw)
 Gain (dB)*46.69963
 Maximum detection range (km)*150.0
 Radar cross section (dB-m⁻²)*0.0
 System losses (dB)
 Receiver noise temperature (dB K)..
 SNR required for detection (dB) ..
 Noise figure (dB)
 Doppler filter bandwidth (hz)*0.2
 Nominal wavelength (m)*0.03
 Modulation type (pulsed,p-dop)....*p-dop

Figure 7-10: Pulsed Doppler scenario: Radar Range equation (inferred terms)

$$P = \frac{(4\pi)^3 R^4 (S/N) (kT) (N_F) (L)}{G^2 \lambda^2 \sigma (T_I)}$$

- INFERRED PARAMETERS

- G: ANTENNA GAIN (FROM ANTENNA SIZE AND RADIO FREQUENCY)
- λ : WAVELENGTH (FROM RADIO FREQUENCY)
- σ : RADAR CROSS SECTION (FROM TARGET FORM INPUTS)
- T_I : INTEGRATION TIME (FROM PULSED-DOPPLER INFERENCE AND PULSE TRAIN DURATION)
- R: MAXIMUM DETECTION RANGE (FROM UNAMBIGUOUS RANGE OF SEARCH DWELL)

- ANALYST SUPPLIED PARAMETERS

- S/N: SIGNAL TO NOISE RATIO REQUIRED FOR DETECTION
- T: RECEIVER NOISE TEMPERATURE
- N_F : NOISE FIGURE
- L: SYSTEM LOSSES

- GOAL PARAMETER

- PEAK POWER (INFERRED FROM DUTY CYCLE, PULSED DOPPLER AND SOLVING RADAR RANGE EQUATION FOR AVERAGE POWER)

Figure 7-11: Pulsed Doppler scenario: Radar range equation used to estimate peak transmitter power

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

Radar range equation parameters

Average power (Kw)*0.162607
 Peak power (Kw)*1.806753
 Gain (dB)*46.69963
 Maximum detection range (km)*150.0
 Radar cross section (dB-m²)*0.0
 System losses (dB) 15.0
 Receiver noise temperature (dB K) .. 24.6
 SNR required for detection (dB) .. 4.0
 Noise figure (dB) 7.0
 Doppler filter bandwidth (hz)*0.2
 Nominal wavelength (m)*0.03
 Modulation type (pulsed,p-dop)....*p-dop

Figure 7-12: Pulsed Doppler scenario: Results of using Radar Range equation to estimate peak transmitter power

Figure 7-12 illustrates the state of the radar range equation form after entering the data and inferencing the peak power. ASTA has used the radar range equation to calculate the average power and the duty cycle obtained from ELINT to convert the result to peak power.

The final phase of the pulsed Doppler scenario consists of an ELINT observation of the radar during initial tracking tests. An additional phase-shift-keyed pulse compression frame is observed performing track after every sixth search dwell. The frame characteristics of the new frame are shown in Figure 7-13. The unambiguous range of the track dwell is less than the range resolution of the search frames facilitating handover from search mode to track mode. The high PRF results in a large unambiguous velocity.

A summary of the resolution and measurement accuracy for each measured parameter for each frame is given in Figure 7-14. The range resolution of the *higher PRF search waveforms* is improved slightly because of the shorter pulse lengths used to maintain constant duty cycle. The effect of the pulse compression in the tracking frame range accuracy is evident. The range-rate resolutions and accuracy of a pulsed Doppler system depend primarily on pulse train duration and frequency which are the same for all frames. Likewise, the angular measurements have the same resolution and accuracy because the frequency is constant from frame to frame. The asymmetry in antenna height and width is reflected in the angular resolution and accuracy.

The measurement accuracy improvement obtained by combining three search frames at the dwell level can be seen in Figure 7-15. Figure 7-16 shows the characteristics of the search track sequence. The target revisit interval is once every 20 frames or 116 milliseconds, for a data rate of 8.6 Hz. The target handling capacity calculated by determining the number of track dwells that can be fit into the search track sequence is $20/2 = 10$.


```

Advanced Information & Decision Systems                                ASTA EXPERIMENTAL SYSTEM
-----
_Help          _Debug modes    _MRS direct    _Interceptor  _Quit ASTA
_System        _Transmitter    _Antenna       _Target       _Range equation
_Sequences     _Dwell types    _Frame types   _Checkpoint   _Match table
-----
Frame Types Menu:  _Add frame type  _Accuracy table  _Return to top menu
_f1 _f2 _f3 _f4
-----
FRAME TYPE:  f4
RF (GHz) ..... 10.0          PRF mode (high/medium/low) .....
PRF (KHz) ..... 109.6        Duty cycle (%) .....*10.96
PRI (us) .....              Unambiguous range (Km) .....*1.368613
Pulse shape ..... rect      Unambiguous velocity (m/s) .....*1644.0
Pulse duration (us) .... 1.0  Pulses per frame .....*548.0
Pulse train duration(ms) 5.0  Waveform energy (joules) .....
Frame duration (ms) .... 5.8  Min range tracking loop bndwth(Hz)*0.869565
Probable oper. mode..... track  Min range tracking accuracy (m) ...
Pulse modulation ..... psk     Max range tracking accuracy (m) ...
Frequency excursion(MHz)      Min SNR to avoid track loss (dB) ..
Chip duration (usec)..... 0.1

```

Figure 7-13: Pulsed Doppler Scenario: Frame characteristics of PSK
pulse-compressed tracking frame f1

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table
/				
Return to frame level (^E)		ACCURACY TABLE		
		FRAME LEVEL		
RESOLUTION	RANGE (Meters)	RANGE RATE (Meters/sec)	AZIMUTH (Degrees)	ELEVATION (Degrees)
(Frame 1)	*1350.0	*4.5	*1.143	*0.9144
(Frame 2)	*1012.5	*4.5		
(Frame 3)	*855.0	*4.5		
(Frame 4)	*22.5	*4.5	*1.143	*0.9144
(Frame 5)				
SINGLE FRAME ACCURACY				
(Frame 1)	*90.0	*0.3	*0.0403174	*0.0322539
(Frame 2)	*67.5	*0.3	*0.0403174	*0.0322539
(Frame 3)	*57.0	*0.3	*0.0403174	*0.0322539
(Frame 4)	*1.5	*0.3	*0.0403174	*0.0322539
(Frame 5)				

Figure 7-14: Pulsed Doppler scenario: Resolution and measurement accuracy
at the frame level

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

Return to dwell level (^E)		ACCURACY TABLE		
		DWELL LEVEL		
SINGLE DWELL ACCURACY	RANGE (Meters)	RANGE RATE (Meters/sec)	AZIMUTH (Degrees)	ELEVATION (Degrees)
(Dwell 1)	*39.201444	*0.1732050	*0.0381234	*0.0304987
(Dwell 2)	*1.0606601	*0.2121320	*0.0466914	*0.0373531
(Dwell 3)				
(Dwell 4)				
(Dwell 5)				

Figure 7-15: Pulsed Doppler scenario: Measurement accuracy at the dwell level

```

Advanced Information & Decision Systems                                ASTA EXPERIMENTAL SYSTEM
-----
_Help          _Debug modes      _MRS direct    _Interceptor   _Quit ASTA
_System        _Transmitter      _Antenna       _Target        _Range equation
_Sequences     _Dwell types      _Frame types   _Checkpoint    _Match table
-----
Sequences Menu:      _Add a sequence                                _Return to top menu
_sl
-----
SEQUENCE:  sl
  Track Dwells in this sequence
  Dwell-type  Dwell-repetitions
1.    d1      6
2.    d2      1
3.
4.
5.
6.
7.
8.
9.
10.
Sequence duration (msec) .....*116.0
Probable operational mode .....*mixed
Sequence length (frame count).....*20
Target revisit interval (msec).....*116.0
Data rate (Hz) .....*8.620689
Probable target handling capacity ..*10
Max range tracking loop bndwdth (Hz)*4.310344
Probable number of target revisits.. 1

```

Figure 7-16: Pulsed Doppler scenario: Characteristics of search track sequence

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

TARGET

Target characteristics

Target type (small).....*small

Cross-section (dB-m sq) 0.0

Maximum velocity (m/sec) 600.0

Maximum acceleration (m/sec²).. 100.0

Figure 7-17: Pulsed Doppler scenario: Target form

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	

_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

INTERCEPTOR

Interceptor characteristics

Maximum range (km)	120.0
Maximum velocity (km/sec)	1.5
Maximum detection range (km)*	171.0
Flyout time (sec)*	80.0
Time from detect to intercept (sec)*	85.0
Reaction time (sec)	5.0

Figure 7-18: Pulsed Doppler scenario: Interceptor form

Figure 7-17 and Figure 7-18 shows the target and interceptor characteristics, respectively. The target assumptions are consistent with a cruise missile. The maximum detection range calculated via engagement analysis (171 km) is consistent with the maximum detection range calculated in an unambiguous range analysis (150 km).

Figure 7-19 gives a summary of the analysis of the radar system. The modulation type (pulsed doppler) and operational mode (search track) have been determined. A summary of the search mode characteristics (unambiguous range and velocity, maximum detection range) and track mode characteristics (tracking loop bandwidths, measurement accuracies, target revisit interval) have been obtained.

7.2 ANALYSIS SCENARIO 2: MECHANICALLY-SCANNED EARLY WARNING RADAR

The radar configuration considered in this analysis scenario consists of a pulsed early warning radar employing a mechanically rotated antenna. When MTI cancellation is employed with a low PRF radar, blind speeds occur at multiples of the PRF, analogous to the ambiguous velocity concept for pulsed doppler radars. In the example, it is assumed that the radar employs a 10% PRI stagger to eliminate these blind speeds. It is shown that the same calculations (PRI harmonic analysis) that are used to estimate the unambiguous velocity of a pulsed doppler radar can be used to estimate the first blind speed of a pulsed, MTI radar.

The observations in this scenario consist of ELINT, PHOTINT, and associated interceptor and target assumptions as before. In this scenario, a measurement of effective radiated power is assumed in order to demonstrate an alternative use of the radar range equation. The same basic radar range equation used in the previous Pulsed Doppler example is employed with the terms properly interpreted to deal with the pulsed case. The analysis objectives in this scenario

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

RADAR SYSTEM

Nationality (us/ussr)
 Operational modes (s,t,mixed)...*mixed
 Modulation type (pulsed,p-dop)...*p-dop
 Unambiguous range (km)*150.0
 Unambiguous velocity (m/s)*1800.0
 Max range track loop bndwidth(Hz)*4.310344827
 Min range track loop bndwidth(Hz)*0.354998513
 Tracking accuracy
 Velocity (km/s)
 Azimuth (deg)*0.009849123
 Elevation (deg)*0.007879298
 prob. target handling capacity .
 Peak ERP (dBwatt)
 Target revisit interval (msec)..*116.0
 Maximum detection range (km)....*150.0

Figure 7-19: Pulsed Doppler scenario: Radar System analysis summary

are to determine the maximum detection range using the radar range equation, to determine the transmitter output tube type, and to estimate the range and angle resolutions. The purposes of this scenario are:

1. To illustrate how the rotating beam characteristics of the radar induce an implicit frame/dwell/sequence description of the radar waveform that was explicitly defined for the electronically scanned scenario.
2. To illustrate the analogy between the use of discrete multiple PRFs to increase the unambiguous range and unambiguous velocity of pulsed doppler radars and the use of pulse stagger to eliminate blind speeds in a conventional pulsed, MTI radar, and to show how the knowledge representation and explanation capability employed in ASTA both clarifies this analogy and simplifies generalization of the computational algorithms.
3. To illustrate the application of the transmitter subsystem rules to a radar system that does not require pulse-to-pulse coherence. ASTA does not currently contain a rule base for obtaining a complete description of the MTI subsystem. It is assumed in this scenario that the analyst hypothesizes noncoherent MTI by entering *noncoherent* on the transmitter form. This represents one possible interpretation of the observed data. Other hypotheses can be pursued by using the *change-hypothesis* mechanism. Our purpose here is to illustrate the differences in the transmitter analysis that result from changing the assumptions of Scenario 1, not to provide a detailed analysis of the MTI subsystem.
4. To illustrate an alternative use of the radar range equation.
5. To illustrate how the *change hypothesis* function in ASTA can be used to efficiently calculate the characteristics of a multiple beam antenna

system.

The waveform employed by the radar consists of a nominal 500 microsecond PRI with a 10% stagger. The stagger sequence consists of alternating long and short PRIs. Figures 7-20 and 7-21 show the waveform data entered by the analyst to describe the basic waveform. A separate frame has been defined for each PRI. The analyst has entered 1 in the pulses per frame slot to indicate that each frame comprises a single pulse.

The pulse train duration is set equal to the pulse duration and the frames duration is set equal to the PRI by the inference engine. Note that the unambiguous range of 75 km is adequate for the assumed maximum interceptor effective range of 30 km, but that the first blind speed (unambiguous velocity) of 60 m/s is inadequate. This situation is typical of the low PRF radar under study in this scenario.

Figure 7-22 illustrates the relationship between the frame and dwell characteristics for the step scanned antenna of scenario 1 and the circularly scanned antenna under consideration in this scenario. In the former case, the frames are pulse trains separated by dead zones where the frames belonging to a common dwell have the same envelope amplitude. In the current case, the circular scanning and beam pattern cause a smooth variation in the envelope amplitude. Each frame consists of a single pulse and the dwell is described as an alternating sequence of frames of type 1 and 2, the total length of which is determined by the 3dB beamwidth of the antenna.

Figure 7-23 shows the dwell form for this radar. The number of repetitions of each frame in the dwell may be either entered by the analyst by consideration of the amplitude variations, as illustrated in Figure 7-21, or inferred from the beam characteristics and measured PRIs, if both ELINT and PHOTINT are

```

Advanced Information & Decision Systems                                ASTA EXPERIMENTAL SYSTEM
-----
_Help          _Debug modes      _MRS direct     _Interceptor    _Quit ASTA
_System        _Transmitter      _Antenna        _Target         _Range equation
_Sequences    _Dwell types       _Frame types    _Checkpoint     _Match table
-----
Frame Types Menu:  _Add frame type  _Accuracy table  _Return to top menu
_f1 _f2
-----
FRAME TYPE:  f1
RF (GHz) ..... 5.0          PRF mode (high/medium/low) .....
PRF (KHz) .....*2.0        Duty cycle (%) .....*0.1
PRI (us) ..... 500.0       Unambiguous range (Km) .....*75.0
Pulse shape ..... rect    Unambiguous velocity (m/s) .....*60.0
Pulse duration (us) ..... 0.5  Pulses per frame ..... 1
Pulse train duration(ms)*0.0005  Waveform energy (joules) .....
Frame duration (ms) .....*0.5  Min range tracking loop bndwdth(Hz)
Probable oper. mode..... search  Min range tracking accuracy (m) ...
Pulse modulation ..... unmod    Max range tracking accuracy (m) ...
Frequency excursion(MHz)         Min SNR to avoid track loss (dB) ..
Chip duration (usec).....

```


Figure 7-20: Pulsed Early-warning scenario: Search frame f1

```


Advanced Information & Decision Systems                                ASTA EXPERIMENTAL SYSTEM
-----
 Help      Debug modes      MRS direct  Interceptor  Quit ASTA
System     Transmitter      Antenna     Target       Range equation
Sequences  Dwell types      Frame types Checkpoint  Match table
-----
Frame Types Menu:  Add frame type  Accuracy table  Return to top menu
_f1 _f2
-----
FRAME TYPE:  f2
RF (GHz) ..... 5.0      PRF mode (high/medium/low) .....
PRF (KHz) ..... *1.818181 Duty cycle (%) ..... *0.090909
PRI (us) ..... *550.0000 Unambiguous range (Km) .....
Pulse shape ..... rect  Unambiguous velocity (m/s) .....
Pulse duration (us) .... 0.5  Pulses per frame ..... 1
Pulse train duration(ms) *0.0005 Waveform energy (joules) .....
Frame duration (ms) ....      Min range tracking loop bndwidth(Hz)
Probable oper. mode..... search Min range tracking accuracy (m) ...
Pulse modulation ..... unmod  Max range tracking accuracy (m) ...
Frequency excursion(MHz)      Min SNR to avoid track loss (dB) ..
Chip duration (usec).....

```

Figure 7-21: Pulsed Early Warning scenario: Search frame f2



STAGGERED-PRF EARLY WARNING RADAR
WITH SMOOTHLY SCANNING ANTENNA



MULTIPLE DISCRETE PRF PULSED-DOPPLER RADAR
WITH ELECTRONICALLY STEP-SCANNED ANTENNA

Figure 7-22: Comparison of frame dwell definitions for (a) circular scan and
(b) step scan

```

Advanced Information & Decision Systems                                ASTA EXPERIMENTAL SYSTEM
-----
_Help          _Debug modes      _MRS direct     _Interceptor    _Quit ASTA
_System        _Transmitter      _Antenna        _Target         _Range equation
_Sequences     _Dwell types      _Frame types    _Checkpoint     _Match table
-----
Dwell Types Menu:  _Add dwell type  _Accuracy table  _Return to top menu
_d1
-----
DWELL TYPE:  d1
Sequence of frames in this dwell
Frame-type Repetitions Phantom-PRF
1.  f1          10
2.  f2          10
3.
4.
5.
6.
7.
8.
9.
10.
11.
Max pulse train duration (msec)*0.0005
Unambiguous range (km).....*825.0
Unambiguous velocity (m/s) ....*600.0
Probable operational mode .....
Min range tracking accuracy (m)
Max range tracking accuracy (m)
Fundamental PRF (KHz).....*20.0
PRF minimum (KHz).....*1.81818181
Ranging PRF (KHz).....*0.18181818
Nominal Frequency (MHz).....*5.0

```

Figure 7-23: Pulsed Early-warning scenario: Search dwell d1


available.

The use of stagger has increased the first blind speed to 600 m/s and concomitantly increased the unambiguous range to 825 km. The latter is a side effect of the designer's explicit purpose of increasing the blind speed. For this reason it is inappropriate to use unambiguous range as an estimate of maximum detection range for a pulsed system of this type. Examination of the radar system form at this stage of the analysis would indicate that no value for maximum detection range has been derived.

Figure 7-24 illustrates how the sequence characteristics of a pulsed system are implicitly defined by the antenna beam characteristics and the scan characteristics of the antenna. In the current case of a circularly scanned antenna, the number of dwells in the sequence is simply the number of 3dB beamwidths in the circular sweep of the antenna. Alternatively, if the scan time of the antenna is known, then the number of dwell repetitions may be calculated from the ratio of scan time to dwell duration.

Figure 7-25 illustrates the sequence form for the pulsed early warning scenario. In this case the sequence duration may be directly related to the time interval between revisits to the same spatial location.

Figure 7-26 shows the PHOTINT data obtained for this radar. Fourier aperture analysis has been used as in the previous scenario to obtain estimate of antenna gain and beamwidths. Figure 7-27 shows the transmitter subsystem inferences. The analyst has entered the assumption that coherent combining techniques are not likely. ASTA concludes that the transmitter oscillator configuration is modulated-final-amplifier (MFA) (also called power oscillator) and that a magnetron output tube is likely. Also the peak output power of the transmitter has been estimated at 16.5 kw, using the measured peak ERP and calculated antenna gain.



SEQUENCE DURATION = SCAN TIME

NUMBER OF DWELLS IN SEQUENCE = #3DB BEAMWIDTHS IN CIRCLE

FRAME DURATION = PULSE REPETITION INTERVAL

NUMBER OF FRAMES IN DWELL = TIME TO SCAN 1 BEAMWIDTH/NOMINAL PRI

Figure 7-24: Defining sequence by antenna beam and scan characteristics


```

Advanced Information & Decision Systems                                ASTA EXPERIMENTAL SYSTEM
-----
_Help          _Debug modes      _MRS direct    _Interceptor  _Quit ASTA
_System        _Transmitter      _Antenna       _Target       _Range equation
_Sequences     _Dwell types      _Frame types   _Checkpoint   _Match table
-----
Sequences Menu:      _Add a sequence                                _Return to top menu
_sl
-----
SEQUENCE:  sl
  Track Dwells in this sequence
  Dwell-type  Dwell-repetitions
1.      dl          257
2.
3.
4.
5.
6.
7.
8.
9.
10.
Sequence duration (msec) .....*2698.5
Probable operational mode .....*search
Sequence length (frame count).....*5140
Target revisit interval (msec).....*2698.5
Data rate (Hz) .....*0.370576
Probable target handling capacity ..
Max range tracking loop bndwdth (Hz)
Probable number of target revisits.. 1

```

Figure 7-25: Pulsed Early-warning scenario: Sequence form

Advanced Information & Decision Systems		ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor _Quit ASTA
_System	_Transmitter	_Antenna	_Target _Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint _Match table

RADAR ANTENNA

Type(planar, horn+refl).....	planar	Beam Characteristics	
Physical Description		beam width - az (deg)	*1.42628
shape(rect,circ).....	rect	beam width - el (deg).....	*3.328
Dimensions		gain (dB)	*40.8909
height (m)	1.5	first sidelobe level (dB) ..	*-31
width (m)	3.5	efficiency	*0.67
radius if horn+r (m) ...		nominal wavelength(m).....	*0.06
Area (m ²)	*5.25	polarization control (yes/no) ..	
aperture			
distribution.....	cos-sq		

Figure 7-26: Pulsed Early-warning scenario: Antenna form

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

RADAR TRANSMITTER

Oscillator configuration (mopa,mfa).....*mfa

Peak power output (KWatt)*165.92075570

Average power output (KWatt)

Nominal frequency (GHz)*5.0

Operating frequency band

Fractional instantaneous bandwidth.....*0.0003544

Coherent? (yes/no)*no

Instantaneous bandwidth (Mhz)*1.772

Maximum duty cycle*0.1

Minimum duty cycle*0.0909090909

Tube type likelihoods:

TWT likelihood

CFA likelihood

Klystron likelihood*unlikely

Magnetron likelihood*likely

Figure 7-27: Pulsed Early-warning scenario: Transmitter subsystem inferences

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

Radar range equation parameters

Average power (Kw)*0.162607
Peak power (Kw)*1.806753
Gain (dB)*46.69963
Maximum detection range (km)*150.0
Radar cross section (dB-m²)*0.0
System losses (dB) 15.0
Receiver noise temperature (dB K) .. 24.6
SNR required for detection (dB) .. 4.0
Noise figure (dB) 7.0
Doppler filter bandwidth (hz)*0.2
Nominal wavelength (m)*0.03
Modulation type (pulsed,p-dop)....*p-dop

Figure 7-28: Pulsed Early-warning scenario: Radar Range equation inferred terms

Table 7-3: Pulsed Early-warning scenario: Method of deriving
radar range parameters

$$R^4 = \frac{PG^2\lambda^2\sigma T_I}{(4\pi)^3(S/N)(kT)N_f L}$$

Inferred parameters:

- P : Peak transmitter power (from pulsed, antenna gain inference, and measured ERP)
- G : Antenna gain (from antenna size and radio frequency)
- λ : Wavelength (from radio frequency)
- σ : Radar cross section (from target form data)
- T_I : Integration time (from pulsed inference and pulse duration)

Analyst-supplied parameters:

- S/N : Signal-to-noise ratio required for detection
- T : Receiver noise temperature
- N_f : Noise figure
- L : System losses

Goal parameter:

- R : Maximum detection range

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	

_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

Radar range equation parameters				
Average power (Kw)*0.150837				
Peak power (Kw)*165.9207				
Gain (dB)*40.89099				
Maximum detection range (km)*43.48470				
Radar cross section (dB-m ²)*0.0				
System losses (dB) 15.0				
Receiver noise temperature (dB K). 24.6				
SNR required for detection (dB) .. 10.0				
Noise figure (dB) 7.0				
Doppler filter bandwidth (hz)*2000.0				
Nominal wavelength (m)*0.06				
Modulation type (pulsed,p-dop)....*pulsed				

Figure 7-29: Pulsed Early-warning scenario: Maximum Detection Range
inference results

```
Advanced Information & Decision Systems          ASTA EXPERIMENTAL SYSTEM
-----
_Help          _Debug modes      _MRS direct    _Interceptor   _Quit ASTA
_System        _Transmitter      _Antenna       _Target        _Range equation
_Sequences     _Dwell types       _Frame types   _Checkpoint    _Match table
-----

TARGET

Target characteristics
Target type (small).....
Cross-section (dB-m sq) ..... 0.0
Maximum velocity (m/sec) ..... 600.0
Maximum acceleration (m/sec^2).. 300.0
```

Figure 7-30: Pulsed Early-warning scenario: Target characteristics

```
Advanced Information & Decision Systems                                ASTA EXPERIMENTAL SYSTEM
-----
_Help          _Debug modes      _MRS direct     _Interceptor    _Quit ASTA
_System        _Transmitter      _Antenna        _Target         _Range equation
_Sequences     _Dwell types       _Frame types    _Checkpoint     _Match table
-----

INTERCEPTOR

Interceptor characteristics
Maximum range (km) ..... 30
Maximum velocity (km/sec) ..... 1.0
Maximum detection range (km) .....*51.0
Flyout time (sec) .....*30.0
Time from detect to intercept (sec)*35.0
Reaction time (sec) ..... 5.0
```

Figure 7-31: Pulsed Early-warning scenario: Interceptor analysis

Figure 7-28 illustrates state of the radar range equation after pressing the fill key. Table 7-3 indicates the method of derivation of each of the inferred parameters. Note that ASTA has used the pulsed inference in determining that peak transmitter power and pulse duration are the appropriate interpretations of P and T_p , respectively. At this point, no independent estimate of maximum detection range is available, so the analyst asserts values of the analyst supplied parameters in Table 7-3. The result of pressing the **infer** key on maximum detection range is illustrated in Figure 7-29. The radar range equation has been solved in the form given in Table 7-3, yielding an estimated maximum detection range of 43 km. Figure 7-30 and Figure 7-31 indicate that the maximum detection range (51 km) obtained by applying engagement analysis to the assumed target and interceptor characteristics is consistent with the value obtained by using the radar range equation (43 km).

Finally, Figure 7-32 illustrates the system form for the radar of scenario 2.

Advanced Information & Decision Systems			ASTA EXPERIMENTAL SYSTEM	
_Help	_Debug modes	_MRS direct	_Interceptor	_Quit ASTA
_System	_Transmitter	_Antenna	_Target	_Range equation
_Sequences	_Dwell types	_Frame types	_Checkpoint	_Match table

RADAR SYSTEM

Nationality (us/ussr)

Operational modes (s,t,mixed)...

Modulation type (pulsed,p-dop)..*pulsed

Unambiguous range (km)*825.0

Unambiguous velocity (m/s)*600.0

Max range track loop bndwdth(Hz)

Min range track loop bndwdth(Hz)*2.129991080

Tracking accuracy

 Velocity (km/s)

 Azimuth (deg)

 Elevation (deg)

prob. target handling capacity .

Peak ERP (dBwatt) 93.0

Target revisit interval (msec)..

Maximum detection range (km)....

Figure 7-32: Pulsed Early-warning scenario: Radar System summary

8. SUMMARY

We have described a research prototype software system named ASTA that employs artificial intelligence technology to interpret radar signals and infer the probable nature of the radar system that produces them. Although ASTA is certainly incomplete in its radar knowledge and still somewhat brittle during operation, this system is capable of yielding analyses of radar signals of high quality, perhaps generally at the level of a well-schooled junior analyst and, for at least some analyses, with a surprisingly high degree of detail and fidelity.

Advances in the state of expert systems research that have occurred in the course of this project include at least the following:

- We have codified some of the knowledge of radar system analysts, including such areas of expertise as signal analysis, antenna and feedline analysis, and knowledge of historical technological developments, and we have embedded that knowledge in an expert system capable of doing a passable job of inferring the design characteristics of the signal source. A handful of other researchers have employed structural models of electronic components, such as integrated circuits or printed circuit boards, to support a limited range of reasoning (e.g. fault diagnosis) about the behavior of those components. Those system typically start with knowledge of both the design principles and system structure, and then attempt to infer likely lines of component behavior given that some fault occurs. Conversely, the ASTA effort represents the first known attempt to structure knowledge about an entire electronic system in terms of its observable operating characteristics, so that its *design principles* can be inferred purely from its behavior and the general principles of operation

of signal-generating systems.

- We have demonstrated an approach to user interface design that incorporates artificial intelligence into a system that is perceived by the user, for the most part, as a spreadsheet program. Forward and backward chaining, remote database update, and a mixed initiative control strategy were smoothly integrated into the spreadsheet metaphor in order to make the system seem as simple and coherent as possible.
- We have identified several classes of conflict that can occur within an expert system. We have further proposed, constructed, and demonstrated techniques that are capable of either resolving the conflict or else detecting the conflict and declaring it to the user, including the capability of providing an English-like explanation of the source of the conflict upon request.
- We have developed a model for the interaction of an expert system with a remote database, such as might be managed by a conventional database management system (DBMS). The remote database is expected to be volatile in schema as well as contents and to model its contents in a fashion that is not trivially compatible with the expert system's representation of the data of mutual interest. Specifically, our approach provides for the interposition of a database access agent between the DBMS and the expert system to perform the necessary mapping between data views and to isolate the expert system from changes in the DBMS's schema. We have also advocated a new class of database access protocols that allow queries to be approximate, rather than completely specific, in nature; and we have proposed a mechanism (the passing of active procedural predicate objects to the DBMS, rather than passive query text objects) that could support such a capability.

- We have developed a viable loosely-coupled multiple-process expert system architecture in which functionality can be separated into distinct "agents" capable of solving distinct tasks (in our case, management of inference and management of user interaction, respectively). This architecture has been implemented and demonstrated in a working research prototype expert system that can solve some relatively realistic analysis problems.

For additional detail and other perspectives on our work, the reader is referred to the Bibliography, which includes several papers that were published based on the technology developed in this project.

It is thus our feeling that the ASTA project has led to several important advances of both a technological and a practical nature. Further work is still required to make the system more robust and "bullet-proof," if it is to be of significant practical value to radar system analysts on a day-to-day basis; this work includes both reimplementation (using faster, more sophisticated software technology and workstation hardware) and substantial extension of the knowledge base. ASTA nonetheless constitutes a relatively effective demonstration of the applicability of AI technology to the radar system analysis problem in particular, and the corresponding but broader class of interactive analytic problem-solving domains as well.

APPENDIX A. BIBLIOGRAPHY


- [1] Adams, T., *et al.* (1984). "A Knowledge-Based System for Analyzing Radar Systems" (invited paper). *Proc. Mil. Microwave Conf.* London, England.
- [2] Cromarty, A., *et al.* (1985). "Distributed Database Considerations in an Expert System for Radar Analysis." In *Expert Data Base Systems*, Benjamin/Cummings.
- [3] Cromarty, A. (1985). "Control of processes by communication over ports as a paradigm for distributed knowledge-based system design." (Submitted for publication.)
- [4] Genesereth, M. (1982). An introduction to MRS for AI experts. Technical Report HPP-80-24, Stanford University Heuristic Programming Project, November 1982.
- [5] Rothnie, J., *et al.* (1980). Introduction to a system for distributed databases (SDD-1). *ACM Transactions on Database Systems* 5(1):1-17.
- [6] Skolnik, M. (1970). *Radar Handbook*. McGraw-Hill.
- [7] Smith, L. (1982). *Intelligence Radar Analysis Handbook, Vol. I, Performance Analysis*. 4-Jan-1982.
- [8] Wilson, G., *et al.* (1983). The Multipurpose Presentation System. *Proc. Ninth International Conference on Very Large Data Bases*. Florence, Italy, October 1983.

- [9] Wilson, G., *et al.* (1984). "AI Assists Analysts in Identifying Soviet Radar Systems." *Defense Systems Review*, January 1984.

APPENDIX B. ASTA KEYBOARD FUNCTIONS

A special keyboard layout was designed for the ASTA user interface. In this appendix, we describe the key names and their functions.

ASTA provides a uniform interface across a variety of terminal types by mapping some of the keys on the terminal keyboard to special ASTA functions. In order to facilitate this canonicalization of the user interface, ASTA incorporates software that sends appropriate escape sequences to the terminal to tell the terminal to emit a specific control character sequence for each of several keys. When pressed, these keys subsequently will produce special character sequences that can be interpreted by ASTA as commands. This gives the user an input interface that approximates a special ASTA device having pushbuttons to invoke ASTA function.



Exit Form	Deep Explain	Fill	Report	Explain	Infer	Enter/Select
--------------	-----------------	------	--------	---------	-------	--------------

Figure B-1: ASTA keypad layout

The special keys are layed out in a standard fashion, as depicted in Figure B-1, plus the standard four "arrow keys" having arrows that point up, down, left, and right. The following list defines the function associated with each ASTA key:

- *Exit form*: The analyst presses this key in order to move to the menu form that is up one level in the menu system. If the analyst returns to this form, it will be in the same state as it was when it was exited.
- *Deep explain*: This key invokes the detailed, recursive explanation facility on the fact at the slot currently pointed to by the cursor.
- *Fill*: Pressing this key induces the ASTA system to attempt to fill in (nearly) all slots visible on the current form. In the current implementation, there are two reasons that a slot may not be filled after this key has been pressed:
 - It may be impossible to calculate or infer a value for some slots with the information currently available.
 - The slot may have been specifically excluded by the ASTA implementors from the list of slots to be filled when this key is pressed. Usually this is done only for slots that are expensive to calculate/infer values for and rarely possible to fill.

In a future release of ASTA, it will be possible for the analyst to customize the user interface by specifying a preference as to which slots should be filled when this key is pressed.

- *Report*: Pressing this key results in the generation of a printable report summarizing the content of all ASTA forms in the current analysis session.

- *Explain*: When this key is pressed, an English-like explanation of the line of reasoning is printed for the value in the slot pointed to by the cursor.
- *Infer*: Pressing this key instructs ASTA's inference engine to attempt to find a value for the value corresponding to the slot pointed to by the cursor. Normally, this is achieved through use of the backwards chaining inference technique. While inference is under way, the ASTA user interface displays a "Working" sign and will not respond to other keys. If it was not possible to infer a value for the current slot, it will be left blank. Note that (nearly) all slots may be filled at once using the *Fill* key.
- *Enter/Select*: This key is used under two circumstances:
 - When the analyst wishes to select choice from a menu, or
 - When the analyst has entered data on the screen and now wishes for those data to be delivered to the inference engine for processing.
- → : Moves the cursor one character to the right. If the cursor is at the right boundary of an input box, it moves the cursor to the next input box; if the cursor is at the end of the last box on a line, the cursor is moved to the first input box on the next line.
- ← : Same, but moves cursor to the left.
- ↑ : Moves the cursor up one input box (or menu item).
- ↓ : Moves the cursor down one input box (or menu item).

END

DTIC

8-86